arXiv:1410.1839v1 [cs.DS] 7 Oct 2014

Rejecting jobs to Minimize Load and Maximum Flow-time

Anamitra Roy Choudhury* Syamantak Das† Naveen Garg† Amit Kumar†

Abstract

Online algorithms are usually analyzed using the notion of competitive ratio which compares the solution obtained by the algorithm to that obtained by an online adversary for the worst possible input sequence. Often this measure turns out to be too pessimistic, and one popular approach especially for scheduling problems has been that of "resource augmentation" which was first proposed by Kalyanasundaram and Pruhs. Although resource augmentation has been very successful in dealing with a variety of objective functions, there are problems for which even a (arbitrary) constant speedup cannot lead to a constant competitive algorithm. In this paper we propose a "rejection model" which requires no resource augmentation but which permits the online algorithm to not serve an epsilon-fraction of the requests.

The problems considered in this paper are in the restricted assignment setting where each job can be assigned only to a subset of machines. For the load balancing problem where the objective is to minimize the maximum load on any machine, we give $O(\log^2 1/\varepsilon)$-competitive algorithm which rejects at most an ε-fraction of the jobs. For the problem of minimizing the maximum weighted flow-time, we give an $O(1/\varepsilon^4)$-competitive algorithm which can reject at most an ε-fraction of the jobs by weight. We also extend this result to a more general setting where the weights of a job for measuring its weighted flow-time and its contribution towards total allowed rejection weight are different. This is useful, for instance, when we consider the objective of minimizing the maximum stretch. We obtain an $O(1/\varepsilon^6)$-competitive algorithm in this case.

Our algorithms are immediate dispatch, though they may not be immediate reject. All these problems have very strong lower bounds in the speed augmentation model.

*IBM Research - India, New Delhi. email: anamchou@in.ibm.com

†Department of Computer Science and Engg., IIT Delhi. email: {sdas, naveen, amitk}@cse.iitd.ac.in

1 Introduction

Online algorithms are usually analyzed using the notion of competitive ratio which compares the solution obtained by the algorithm to that obtained by an offline adversary for the worst possible input sequence. Researchers have tried to address the criticism of this measure being too pessimistic by either limiting the power of the adversary – oblivious adversary, stochastic adversary – or giving more power to the online algorithm – lookahead, additional resources, etc. One popular approach especially for scheduling problems has been that of "resource augmentation" and was first proposed by Kalyanasundaram and Pruhs [29]. In this model the machines of the online algorithm have more speed than those of the offline algorithm. Many scheduling problems for which no constant competitive online algorithm is possible now have such algorithms in this resource augmentation model. The success of the speed augmentation model lies in the fact that many natural algorithms can be analysed in this framework.

In this paper, we propose a "rejection model" in which there is no resource augmentation, but we allow the online algorithm to not serve an ε-fraction of requests. There are two principal reasons for considering this model: (i) although resource augmentation has been very successful in dealing with a variety of objective functions, there are problems for which even a (arbitrary) constant speedup cannot lead to a constant competitive algorithm – we consider two such problems in this paper, and (ii) this might be a natural assumption in many settings where job rejection is part of the service provided by a system (e.g., "Server busy: Please try again later" message we often see when accessing popular websites).

For most scheduling problems, an algorithm in the resource augmentation model can be "simulated" in this rejection model by roughly letting each machine drop every $(1/\varepsilon)^{th}$ job assigned to it. However, the rejection model is much more powerful than the resource augmentation model since we are not restricted to drop an ε-fraction of jobs assigned to each machine, i.e., we could drop many more jobs assigned to one machine as compared to another as long as the overall number of rejected jobs stays within ε-fraction of the total number of jobs. We demonstrate this by considering two classical problems – load balancing and maximum (weighted) flow time. For both these problems we give constant competitive algorithms in the rejection model while no such algorithms are possible in the resource augmentation model. This is the key contribution of this paper.

The problems considered in this paper are in the restricted assignment setting where each job can be assigned only to a subset of machines. All our algorithms are pre-emptive, immediate dispatch – a job is assigned to a machine as soon as it is released and non-migratory – a job is processed only on the machine to which it is assigned. While ideally we would also like to make rejection decisions immediately when the job is released, we also allow for the job being rejected while it is waiting in the queue. We justify this by showing that no online algorithm with immediate dispatch and immediate rejection can be constant competitive for the load balancing problem (and hence, the maximum flow-time problem).

Our Results: For the load balancing problem(`LoadBalancing`) where the objective is to minimize the maximum load on any machine, we give an $O(\log\left(\frac{1}{\varepsilon}\right))$ competitive immediate rejection algorithm when all jobs have unit processing time and the online algorithm can reject an ε-fraction of the jobs. For general processing times our algorithm is not immediate reject and is $O(\log^2\left(\frac{1}{\varepsilon}\right))$-competitive. We show that one cannot get a constant competitive algorithm if we require immediate rejection. Note that there is a $\Omega(\log m)$ lower bound on the competitive ratio of any online algorithm for this problem where m is the number of machines [10]. Further, making the machines an ε-fraction faster has no significant impact on this lower bound.

For the maximum flow-time problem Anand et. al [4] show that no immediate dispatch algorithm can be constant competitive even when the jobs are unit length and we allow resource augmentation. For this setting of unit jobs we show an immediate reject algorithm which is $O(1/\varepsilon)$-competitive and rejects at most an ε fraction of the jobs. When jobs have weights(`WtdMaxFlowTime`), the objective is to minimize the maximum weighted flow-time of a job. For this setting Anand et. al [4] show that no online algorithm can be constant competitive even when we allow non-immediate dispatch and resource augmentation. Our algorithm for this setting is immediate dispatch but not immediate reject and is allowed to reject jobs of total weight at most an ε-times the total weight of all jobs and has a competitive ratio of $O(1/\varepsilon^4)$. We also show that it is not possible to get better than $O(1/\varepsilon)$-competitive algorithm in this model, and that one cannot get a good competitive algorithm if we are required to perform immediate reject and immediate dispatch.

We further generalize our result to a setting where each job has two kinds of weight – rejection weight and flow-time weight. The weighted flow-time of the job is defined as its flow-time times its flow-time weight, and the goal, as before, is to minimize the maximum weighted flow-time of a job. However, the total rejection weight of the jobs which get rejected should be at most ε fraction of the total rejection weight of all the jobs. We obtain an $O(1/\varepsilon^6)$-competitive algorithm for this problem. The problem of minimizing maximum stretch is a special case of this setting in the rejection model. Here, the rejection weights are all unit while the flow-time weights are the inverse of processing sizes.

2 Related Work

Load Balancing. Graham [26] considered this problem in the context of identical machines and showed that the simple greedy heuristic of assigning the next task to the least loaded machine is 2-competitive (see also the survey by Azar [8]). Albers [1] improved the competitive ratio to 1.923 and also showed a lower bound of 1.852 on the competitive ratio of any deterministic online algorithm, while Albers et al. [2] shows improved bounds in the special case where the online algorithm knows the sum of job sizes at any point. For related machines model, Berman et al. [14] gave constant competitive algorithms. However, the problem becomes significantly harder in the unrelated machines model, where a job can have different processing time on different machines. Azar et al. [10] considered the problem in the restricted assignment setting, and gave an $O(\log m)$-competitive algorithm for load balancing(m being the number of machines). They also complemented this result by proving lower bound of $\Omega(\log m)$ for any deterministic and $\Omega(\ln m)$ for any randomized online algorithm under the restricted assignment model. Buchbinder et al.[15] gave an alternative, more general upper bound on the load on any prefix of the most loaded machines. For the unrelated machines setting, Aspnes et al. [6] gave an $O(\log m)$-competitive algorithm. There has been some work on resource augmentation in this setting. Azar et al.[9] showed a competitive ratio of $1 + 1/2^{\frac{n}{m}(1-o(1))}$ when the online algorithm is allowed to use n identical machines while the offline optimal is restricted to $m < n$ identical machines.

Flow-time minimization. There has been considerable work on scheduling with the objective of minimizing a suitable norm of the flow-time of jobs. For the objective of average flow-time of jobs, a logarithmic competitive algorithm in the identical machines setting is known [30, 7]. Garg and Kumar [23] extended this result to the related machines setting. Garg and Kumar [24] showed that the problem becomes considerably harder in the restricted assignment setting and no online algorithm with bounded competitive ratio is possible. Bansal and Pruhs [12] showed that the competitive ratio can be as high as $\Omega(n^c)$ for the problem of minimizing ℓ_p (for any $1 < p < \infty$)

norm, where n is the number of jobs, even for a single machine. For minimizing the maximum flow-time in the identical machines model, Ambühl and Mastrolilli [3] gave a simple 2-competitive algorithm. However, Anand et al. [4] showed that the competitive ratio of any online algorithm for the restricted assignment setting is as high as $\Omega(m)$, where m is the number of machines.

One approach for circumventing such strong lower bounds has been speed augmentation, where we allow each machine in the online algorithm ε-fraction more speed than the corresponding machine in the offline algorithm. This model was first proposed by Kalyanasundaram and Pruhs [29] who used it to get an $O(1/\varepsilon)$-competitive algorithm for minimizing total flow time on a single machine in the non clairvoyant setting. Bansal and Pruhs [12] proved that several natural scheduling algorithms are $O(1/\varepsilon)$-competitive algorithm for minimizing ℓ_p norm (for any $1 < p < \infty$) of flow-time of jobs in the single machine setting. Golovin et al. [25] extended this result to parallel machines setting. Chekuri et al. [19] showed that the immediate dispatch algorithm of Avrahami and Azar [7] is also $O(1/\varepsilon)$-competitive for all ℓ_p norms ($p \geq 1$).

In the more general setting of unrelated machines with speed augmentation, Chadha et al. [16] gave an $O(1/\varepsilon^2)$-competitive algorithm for minimizing the sum of flow-time of jobs, which was improved and extended to the case of ℓ_p norm of flow-time by Im and Moseley [28] and Anand et al. [5]. However, the competitive ratio remained a function of p, and in fact, Anand et al. [4] showed that one cannot obtain competitive ratio better than $\Omega\left(\frac{p}{\varepsilon^{1-O(1/p)}}\right)$ even in the restricted assignment model.

For the objective of minimizing the maximum (unweighted) flow time on unrelated machines under the speed augmentation model, Anand et al. [4] gave a $O(1/\varepsilon)$-competitive algorithm; however their algorithm is not an immediate dispatch algorithm. In fact, Azar et al. [10] showed that any immediate dispatch algorithm for minimizing maximum flow time will have a competitive ratio can be as high as $\Omega(\log m)$ in the restricted assignment setting even with constant speed augmentation. In the maximum weighted flow-time case, this lower bound holds even if we allow non-immediate dispatch [4]. Chekuri et al. [21, 20] considered the problem of minimizing the maximum delay factor of a job in the parallel machines setting, where jobs come with deadlines and the delay factor of a job is the ratio of its flow-time to the difference between the deadline and its release date. This problem is equivalent to minimizing the maximum weighted flow time, where the weight of a job j can be seen as $(d_j - r_j)^{-1}$. They gave a $(1+\varepsilon)$-speed $O(1/\varepsilon)$-competitive algorithm.

Scheduling with Rejection. There has been considerable work on online scheduling with job rejections in the prize collecting setting. Here, each job comes with a specified penalty which is to be incurred in case it is not scheduled. The goal is to minimize the sum of a suitable objective function of the completion times of the jobs (which are not rejected) and the total penalty cost of rejected jobs. Bartal et al. [13] considered the problem of minimizing makespan in the identical machines model with penalties and gave a 2.618-competitive algorithm. Epstein et al. [22] extended this work to the problem of minimizing makespan on two related machines. Bansal et al. [11] considered the online problem of minimizing total flow-time of jobs and total idle time on a single machine along with uniform rejection penalty for all jobs. They gave 2-competitive algorithms for both the objectives. For the case of arbitrary penalties and average weighted flow-time, they showed strong lower bound of $\Omega(\max(n^{\frac{1}{4}}, C^{\frac{1}{2}}))$ on the competitive ratio of any randomized online algorithm. They complemented this with an $O(\frac{1}{\varepsilon}(\log W + \log C)^2)$-competitive algorithm with $(1+\varepsilon)$-speed augmentation, where W is the ratio between maximum and minimum weights and C is the ratio between maximum and minimum penalties. These results were extended to non-clairvoyant settings (i.e., the size of a job is not known till the time it finishes processing) in the identical machines setting by Chan et al. [17] – they considered an objective function which also

had an energy term, and gave constant competitive $2(1+\varepsilon)$-speed algorithm. Minimizing average flow-time and completion time with rejections has also been studied in the offline context [27].

Charikar and Khuller [18] studied the online problem of minimizing maximum flow-time in more general context of broadcast scheduling, where both the online algorithm and the offline optimum are allowed to ignore a fixed fraction of the jobs. They showed that no randomized online algorithm can be constant competitve, while the offline problem admits a 5-approximation.

3 Problem Statement

We formally define the problems considered in this paper. We consider the online problem of scheduling in the restricted assignment setting. The input instance specifies a small enough positive parameter ε. We have a set of m machines, and jobs arrive in an online manner. A job j arrives at time r_j, and can only be scheduled on a subset S_j of machines. Further, it specifies a processing requirement (or size) of p_j units.

In the `LoadBalancing` problem, a solution needs to dispatch each job to a machine. It can also choose to reject a job (either when it arrives or after a job has been dispatched to a machine). However, for any time t, the total number of jobs rejected by the algorithm till time t must be at most ε fraction of the number of jobs that have arrived till time t. The *load* assigned to a machine at any point of time t is defined as the total processing requirement of jobs which have been dispatched to it till time t – note that this does not count jobs which get rejected by time t, but it counts jobs which will get rejected after time t. The goal of the scheduling algorithm is to minimize (for all time t) the maximum load assigned to any machine.

In the `WtdMaxFlowTime` problem, an input instance can be described as above. Further, we also have a weight w_j associated with each job j. A solution needs to process a job j on one of the machines in S_j for p_j amount of time. Note that any machine can perform 1 unit of processing in unit time, and we allow jobs to be pre-empted. However, migration is not allowed – a job dispatched to a machine i must be completed on i or rejected. The flow-time of a job is defined as the difference between its completion time C_j and its release time r_j. The goal is to minimize the maximum weighted flow-time of a job, i.e., $\max_j w_j \cdot (C_j - r_j)$. As before, till any time t, our algorithm is allowed to reject ε-fraction of jobs which have arrived till this time. This rejection can either happen on arrival of a job or later.

In the `GenWtdMaxFlowTime` problem, the setting is same as above, but the weights for flow-time and the weight for rejection are different. In other words, instead of having one weight w_j, a job j has two weights associated with it, the rejection-weight w_j^r and flow-time-weight w_j^f. Again the objective is to minimize the maximum over all jobs j of $w_j^f F_j$, where F_j denotes the flow-time of job j in a schedule; and we are allowed to reject jobs of total rejection-weight at most ε times the total rejection-weight of all the jobs. Note that in the case of `WtdMaxFlowTime` problem, w_j^r happens to be same as w_j^f.

Our algorithms for `MaxFlowTime` and `WtdMaxFlowTime` (and `GenWtdMaxFlowTime`) satisfy the immediate dispatch property – when a job j arrives, it is dispatched to a machine at time r_j. Note that it may still get rejected later.

We now give an outline of rest of the paper. In Section 4, we give a brief outline of our algorithms and the ideas involved. In Section 5, we give an algorithm for the `LoadBalancing` problem. In Section 6, we describe our algorithm for the `WtdMaxFlowTime` problem. To illustrate some of the ideas, we first consider the special case of unit size and unit weight in Section 6.1. Although one

does not need to invoke an LP relaxation here, we give a more detailed explanation in this section so that similar ideas in later section can be clearer. In Section 6.2, we describe the algorithm for the general case. The algorithm is split in two parts – $\mathcal{A}$ and $\mathcal{B}$. We first describe these two algorithms and give their analysis in Section 6.2.1. In Section 7, we show how to extend our results to the `GenWtdMaxFlowTime` problem. Finally, in Section 8, we give lower bounds on competitive ratios for the problems considered in this paper.

4 Our Techniques

In both the `LoadBalancing` and the `WtdMaxFlowTime` problems, we assume that we know the value of the offline optimum solution, denoted by $T^\star$. Removing this assumption requires standard ideas in online algorithms (details are given in Section 5.3 and Section 6.3).

In the `LoadBalancing` problem, we first consider the special case when all jobs are of unit size. Here, the algorithm is quite natural – when a job arrives, it is dispatched to the least loaded machine (among the ones it can be dispatched to). However, we do not allow the load on a machine to exceed $\alpha T^\star$ where $\alpha = \log\left(\frac{1}{\varepsilon}\right)$ – if we dispatch a job to a machine on which the load is already at this limit, we reject this job. Clearly, the algorithm is $O(\log\left(\frac{1}{\varepsilon}\right))$-competitive. To bound the number of rejected jobs, we use the following argument. Let m_l denote the number of machines where the algorithm dispatches at least $lT^\star$ jobs (l is some integer between 1 and $\log\left(\frac{1}{\varepsilon}\right)$). A naive argument will show that the total number of jobs is at least $m_l \cdot lT^\star$, but a slightly more careful argument can show that the number of jobs in the system is at least $m_l \cdot 2^l T^\star$. A rough argument goes as follows: let M_l be the set of machines where the algorithm assigns at least $lT^\star$ jobs. Consider the jobs dispatched to M_l which are at level $(l-2)T^\star$ or above (level of a job is its position in the queue of the corresponding machine). There will be $2|M_l|T^\star$ such jobs, and so there must be a set of $2|M_l|$ machines where such jobs can be processed (otherwise optimum value cannot be $T^\star$). Each of these machines must have at least $(l-2)T^\star$ jobs assigned to them (because of the job dispatch policy). Thus, the total number of jobs is at least $2(l-2)|M_l|T^\star$. Continuing this argument gives the exponential scaling with l. Now, if the number of rejected jobs is large, then there must be many machines where the load is up to the maximum limit, which means that there must be a huge (exponential in α times number of such machines) number of jobs in the system.

For the more general case of the `LoadBalancing` problem when job sizes are arbitrary, we can reduce the problem to the case of unit size jobs. Assume wlog that all job sizes are powers of 2. For each value of index j, we run the above algorithm independently for jobs of size 2^j. This will ensure that for any fixed j, the total load (i.e., processing time) of jobs of size 2^j assigned to a machine is at most $\alpha T^\star$. In case this limit is reached on a machine for many values of the index j, one can show that we can reject many (large) jobs. Note that this algorithm may reject a job after it gets dispatched to a machine. We show that this is unavoidable (details in Theorem 8.2).

In the `WtdMaxFlowTime` problem, each job j has a weight w_j and the goal is to minimize the maximum over all jobs j, of $w_j F_j$, where F_j is the flow-time of j. In this problem, the ideas are more subtle and technically involved. In fact, we show that unlike the `LoadBalancing` problem, one cannot obtain better than $O(1/\varepsilon)$-competitive algorithm (details in Theorem 8.3). To give intuition about our algorithm and analysis techniques, we consider some special cases:

Unit size and unit weight: Suppose all jobs are of size 1 and weight 1. In this case a natural algorithm is as follows: each machine maintains a queue of jobs assigned to it. When a new job arrives, it is dispatched to the machine with the smallest queue (among the machines to which it can be processed on). However, if all such machines have at least $T^\star/\varepsilon$ jobs in their queue, we

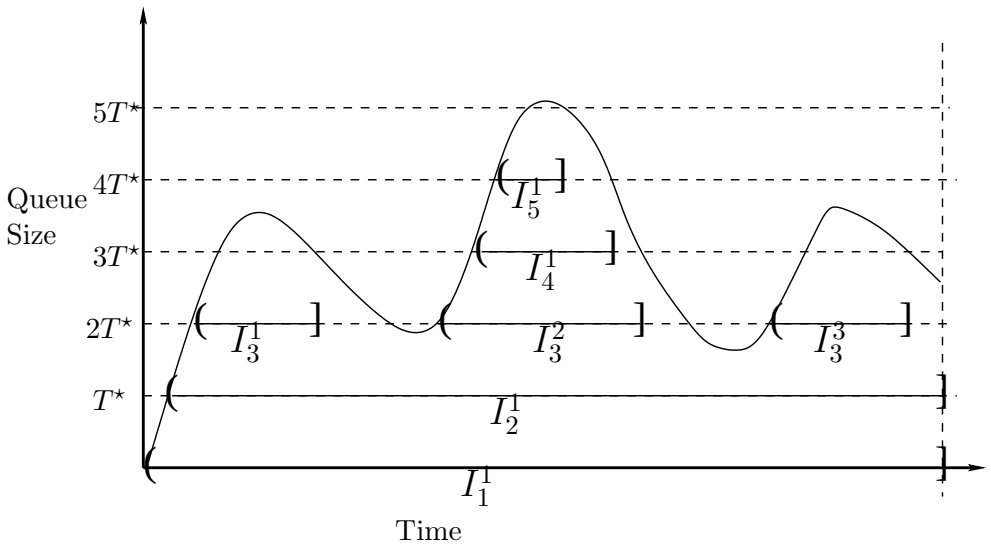

Figure 1: Example showing intervals $\mathcal{I}^{(i,l)}$ for a machine i. Here $\mathcal{I}^{(i,1)} = \{I_1^1\}, \mathcal{I}^{(i,2)} = \{I_2^1\}, \mathcal{I}^{(i,3)} = \{I_3^1, I_3^2, I_3^3\}, \mathcal{I}^{(i,4)} = \{I_4^1\}, \mathcal{I}^{(i,5)} = \{I_5^1\}$

reject this job. The analysis is somewhat trickier than the corresponding case for `LoadBalancing`. The reason is as follows: suppose there is a job j which can go on two machines i_1 and i_2, and both have the same queue size when j is released. Suppose we dispatch j to i_1. But in future, i_1 will continue to get more jobs assigned to it (because perhaps these jobs could only get processed on i_1), whereas i_2's queue will decrease with time. In the `LoadBalancing` problem, queues do not dissipate with time, and so the proof gets much simpler.

Our proof idea is as follows. For each machine i and parameter l (l varies between 1 and $1/\varepsilon$), we define a set of disjoint intervals $\mathcal{I}^{(i,l)}$. An interval I in this set is a minimal interval such that the queue size in i at the left end-point is $(l-1)T^\star$ and that at the right end-point is $lT^\star$ (see Figure 1). For different values of l and fixed i, these intervals form a laminar family. Let m_l be the number of intervals in $\cup_i \mathcal{I}^{(i,l)}$. Suppose the algorithm rejects $kT^\star$ jobs. Then it is not difficult to show that

$$m_{l-1}T^\star + L_{l-1} \geq m_l T^\star + L_l + kT^\star, \tag{1}$$

where L_l denotes the total length of all the intervals in $\cup_i \mathcal{I}^{(i,l)}$. Indeed, the RHS above denotes a set of jobs which get released during an interval in $\cup_i \mathcal{I}^{(i,l)}$, and each such interval can only take $T^\star$ jobs beyond its own length. Summing the above for all l between 1 and $1/\varepsilon$ shows that the number of jobs is at least $kT^\star/\varepsilon$, and so the algorithm rejects only ε fraction of the jobs. The formal proof is given in Section 6.1. Although one can simply formalize the argument given above, we give a somewhat longer proof which sets dual variables for a natural LP relaxation of this problem. This proof generalizes to the more general case.

In more general settings, we split the algorithm in two parts: the first part ensures that the queues on each machine are bounded, and the second part uses this property to ensure that all jobs which do not get rejected finish within the required time. We describe details of the first part below.

Unit weight and arbitrary size For sake of simplicity, assume that job sizes are 1 or 2. Each

machine i maintains two queues: $Q_{i,1}$ and $Q_{i,2}$ for the two job sizes respectively. When a job of size p arrives, it goes to that machine i for which $Q_{i,p}$ has the least load (load of a queue is the total remaining processing time of jobs in it). Again, if the load on all such queues is more than $\alpha T^\star$, we reject the job – here α is a parameter which is $O(1/\varepsilon)$. Thus, the algorithm ensures that all queue sizes remain bounded. The non-triviality lies in figuring out which job a machine chooses to process. Given the above dispatch rule, popular and simple heuristics like always processing the shortest sized job or the job with shortest remaining processing time will not guarantee bounded rejection ratio: Below we provide an example on how this fails.

Assume $\varepsilon = 1/4$, i.e., we are allowed to reject at most $\frac{1}{4}$th of the jobs; and the maximum queue length of any job size which our online algorithm can afford on any machine is $4T^\star$. In this example, $T^\star$ will be 2. Suppose there are 8 machines m_1, m_2, $\cdots$, m_8; also let the jobs can be of two sizes 1 and 2. At time $t = 0$, seven jobs each of size 2 and one of size 1 arrive; the online algorithm sees the jobs in the following order: the first four jobs of size 2 can be processed on any machine, so our algorithm dispatches them to machines, say m_1, m_2, $\cdots$, m_4; the fifth (sixth) job of size 2 can be processed on m_1, m_2 (respectively m_3, m_4), so our algorithm dispatches them to machines m_1 and m_3 respectively; the seventh job of size 2 can be processed on m_1 only. The eighth job of size 1 can be processed on any machine, suppose our algorithm dispatches it to m_1. Thus at end of time $t = 0$, the load on machine m_1 for size 2 jobs and size 1 jobs are 6 and 1 respectively. Since our processing rule picks the shortest sized (or the shortest remaining processing timed) job, the 1 sized job will be processed on machine m_1 at this time step. Now suppose for every time steps $t = 1, 2, \cdots$, a job of size 1 arrives which can be processed on any machine, and for every alternate time steps $t = 2, 4, 6, \cdots$, a job of size 2 arrives which can be processed only on machine m_1. Suppose our algorithm dispatches the size 1 jobs to m_1 (since the algorithm will always prefer processing size 1 jobs over size 2 jobs, the queue size of size 1 jobs on all machines at the end of a time step will be zero, and hence the algorithm may dispatch the size 1 job arriving in the next time step to m_1). The algorithm will start rejecting all the size 2 jobs arriving in time $4, 6, 8, \cdots$. Thus the fraction of the rejected jobs will be close to 1/2.

Note that the optimum solution of the scenario will always dispatch the size 2 jobs of time steps $t = 2, 4, 6, \cdots$ to m_1 and the size 1 jobs to some other machine. Thus the optimal solution will have a maximum flow time of 2.

The processing policy which our algorithm employs is as follows. For a machine i, let $\texttt{load}_{i,1}(t)$ and $\texttt{load}_{i,2}(t)$ denote the load in the two queues at time t. At time t, the machine processes job from the queue for which $\texttt{load}_{i,p}(t)/p$ is largest – this quantity is roughly equal to the number of jobs in the corresponding queue. Let us see why this strategy works. First we consider jobs of size 1. Note that the queue for size 1 jobs can contain up to $\alpha T^\star$ jobs, while that for size 2 jobs can only contain $\alpha T^\star/2$ jobs. As in the above case, we can define the intervals $\mathcal{I}^{(i,l)}$ for the size 1 queues, and write down inequalities (1) as long as $l \geq \alpha/2$ (because as long as the load on size 1 queue is more than $\alpha T^\star/2$, the machine will not give preference to a size 2 job, and so we can pretend that there are size 1 jobs only). This suffices to bound the number of size 1 jobs which get rejected in terms of the total number of size 1 jobs that has arrived.

When we need to bound the number of size 2 rejected jobs, we need to define the intervals more carefully. For a parameter l, $\mathcal{I}^{(i,l)}$ consists of those minimal intervals I where the load of $Q_{i,1}$ is at most $(l-1)T^\star/2$ *and* that of $Q_{i,2}$ is at most $(l-1)T^\star$ at the left end point of I, whereas the load of $Q_{i,1}$ is at least $lT^\star/2$ *or* that of $Q_{i,2}$ is at least $lT^\star$ at the right end-point. Using these intervals, one can again write down inequalities similar to (1) involving both size 1 and size 2 jobs. Using these inequalities, one can prove that the total number of rejected size 2 jobs is at most the total

number of size 2 jobs and *half* the total number of size 1 jobs.

More generally, suppose jobs sizes are powers of 2. Then each machine i maintains a separate queue $Q_{i,k}$ for jobs of size 2^k, and jobs are dispatched to machines according to load on corresponding queues only. At any time t, a machine i processes job from the queue $Q_{i,k}$ for which the load divided by 2^k is highest. Another way of thinking about this rule, which forms the intuition in more general case, is that we prioritize the queues on a machine based on how full they are. Suppose the queue $Q_{i,k}$ is full to a fraction of f_k (i.e., load on it is equal to $f_k \cdot \alpha T^\star$). Then, we prefer the queue for which $f_k/2^k$ is largest.

Arbitrary weight and unit size Again assume that all weights are powers of 2. Each machine i maintains queues $Q_{i,w}$ for jobs with weight 2^w, and as before, jobs are dispatched based on their corresponding queue. A machine i prioritizes these queues in the decreasing order of w. Drawing analogy from the above case, if f_w denotes the fraction to which the queue $Q_{i,w}$ is full (note that a queue for jobs of weight 2^w is full, if the total number of jobs in it is $\alpha T^\star/2^w$, where α is $O(1/\varepsilon)$), then machine i picks the next job from the queue for which $f_w \cdot 2^w$ is highest.

General Case For simplicity, assume all jobs sizes and weights are powers of 2. For each machine i, we have queues $Q_{i,w,p}$ for jobs of size 2^p and weight 2^w. However, it is unclear how to prioritize these queues because we can have two queues Q_{i,w_1,p_1} and Q_{i,w_2,p_2} for which $w_1 < w_2$ but $p_1 < p_2$. Suppose the queues Q_{i,w_1,p_1} and Q_{i,w_2,p_2} are full to the extent of f_1 and f_2 respectively. We know fraction f_1 for Q_{i,w_1,p_1} corresponds to fraction $f_1/2^{p_1-p_2}$ for Q_{i,w_1,p_2} (using the argument above for unit weight and arbitrary size), and the fraction $f_1/2^{p_1-p_2}$ for Q_{i,w_1,p_2} corresponds to fraction $f_1 \cdot \frac{2^{w_1-w_2}}{2^{p_1-p_2}}$ for Q_{i,w_2,p_2} (using the argument for unit size and arbitrary weight). Thus, machine i prefers a job from Q_{i,w_2,p_2} over Q_{i,w_1,p_1} iff $f_2 \geq f_1 \cdot \frac{2^{w_1-w_2}}{2^{p_1-p_2}}$. Our algorithm is based on this rule.

The proof that all queues remain bounded again relies on carefully defining a set of intervals, and corresponding dual variables for an LP relaxation. However there are subtle and non-trivial details. While defining the intervals for a particular type of jobs (say of size 2^p and weight 2^w), it may happen that the algorithm processes other kinds of jobs during these intervals: (i) if these jobs are of very high density, then we are completing lot of weight in small time, and so this should be somehow beneficial, (ii) if there are jobs whose density is close to the density of such jobs, then we cannot handle these directly; in fact, we need to go back and define these intervals for a *group* of jobs of similar density, and (iii) if these jobs have low density, then we need to somehow prove that either they have very high weight and so need to finish soon in any solution, or else such jobs cannot be processed during such intervals.

This completes our informal description of the first part of our algorithm where we ensure that queues remain of bounded size. The second part of the algorithm makes sure that a job of weight w_j does not starve for too long (much longer than a constant times $T^\star/w_j$). To accomplish this, it runs the algorithm above in background. If the first algorithm tries to process a job j at some time t on a machine, the second algorithm processes j as well unless there is a job of much higher density in the queue of machine i. If the latter case happens, it processes such a job. A job which waits for too long (compared to $T^\star/w_j$) gets rejected. We prove that the weight of such jobs is small. The difficulty in the analysis arises for the following reasons: (i) the algorithm $\mathcal{A}$ ensures that the size of each of the queues on a machine remains within a limit, whereas we would like the *total* size of the queues to remain close to this limit; to ensure this, we need to prove that many of these jobs will get rejected, (ii) we may reject a job after processing it for a while – the time used for processing such jobs gets wasted, and so we need to ensure that this remains very small.

Finally, we consider the more general problem `GenWtdMaxFlowTime` where a job j has two weights – w_j^f and w_j^r. Recall that the weight w_j^f is used for computing weighted flow-time, whereas

we do not want to reject jobs whose total w_j^r weight is high. This is a generalised version of the problem of minimizing the maximum stretch. It turns out that the first part of the algorithm described above carries over without much change to this problem as well. However, the second part of the algorithm requires some subtle changes – in particular, we cannot bound the amount of time in which a machine processes jobs which eventually get rejected, and so the above ideas do not apply directly.

5 Algorithm for the `LoadBalancing` problem

In this section, we describe our algorithm for the `LoadBalancing` problem. We first consider a special case when all jobs have size 1. This will illustrate the main ideas involved, and then we shall extend it to the case of general processing time. We shall assume that we know the optimal value, denoted by $T^\star$. We shall show how to get rid of this assumption later.

5.1 Unit size jobs

Our algorithm $\mathcal{A}$ is greedy. We pick a threshold α (which depends on ε) to be defined later. Each machine maintains the current *load* on it, i.e., the number of jobs assigned to it so far (not counting the jobs which have been rejected). Let $\texttt{load}_i(t)$ be the load on machine i at time t. When a job j arrives, it is dispatched to the machine in S_j with the least load, provided the load on this machine is less than $\alpha T^\star$. If the load on this machine happens to be greater than $\alpha T^\star$, we reject this job.

Clearly, our algorithm ensures that we do not exceed the load on a machine by more than $\alpha T^\star$. We need to prove that we will not reject more than ε fraction of the jobs. Suppose this fact is not true, and let n be the first time when we reject more than ε fraction of the jobs – call these rejected jobs J^R.

We divide the machines into several groups. Let M_α denote the set of machines i for which $\texttt{load}_i(n) = \alpha T^\star$. For an integer s, $0 \leq s < \alpha$, let M_s be the set of machines i for which $\texttt{load}_i(n) \geq s \cdot T^\star$. Note that $M_s \subseteq M_{s-1}$. Let m_s denote $|M_s|$. Now we show that m_s increases exponentially as we decrease s.

Consider a job j which arrives at time t. Suppose it is dispatched to machine i. Define $\texttt{level}(j)$ as $\texttt{load}_i(t)$ (not counting j) – we can think of this as the position of j in the queue of machine i. For an index u, let $J(u)$ be the set of jobs whose level is at least u.

Claim 5.1. *For any s, $0 \leq s < \alpha$, $m_s \geq 2^{\alpha-1-s} \cdot m_\alpha$.*

Proof. For any job $j \in J(sT^\star)$, $S_j \subseteq M_s$. Indeed, if a machine $i \notin M_s$, then $\texttt{load}_i(n) < sT^\star$. Further if $i \in S_j$, then j had the option of getting dispatched to i, and then $\texttt{level}(j)$ should be less than $sT^\star$, a contradiction. Therefore, we get the following lower bound on $T^\star$:

$$T^\star \geq \frac{|J(sT^\star)|}{m_s}. \tag{2}$$

Now, it is easy to see that $|J(sT^\star)|$ is at least $T^\star \cdot (m_{s+1} + \ldots + m_\alpha)$. Substituting this in the lower bound above, we get $m_s \geq m_{s+1} + \ldots + m_\alpha$. □

Corollary 5.2. *The number of jobs is at least $(2^\alpha - 1) \cdot T^\star \cdot m_\alpha$.*

Proof. Using the lemma above, the number of jobs is at least

$$(m_0 + m_1 + \ldots + m_{\alpha-1})T^\star \geq (2^\alpha - 1) \cdot m_\alpha T^\star.$$

. □

Corollary 5.2 implies that $|J^R|$ is more than $\varepsilon \cdot (2^\alpha - 1) \cdot T^\star m_0 > T^\star m_0$, if we pick α to be $\log\left(\frac{1}{\varepsilon}\right) + 2$. Now, observe that for any job $j \in J^R$, $S_j \subseteq M_0$. Indeed, the fact this job is rejected means that all the machines in S_j were loaded to the maximum capacity. Now we get a contradiction – we have a set of jobs J^R of cardinality greater than $T^\star m_\alpha$ which can be scheduled only on a subset of m_α machines. Hence, the optimal value must be larger than $T^\star$.

Remarks: If instead of unit size jobs, the processing times of jobs were in the range $[1, 2]$, the above analysis would still apply provided we pick α to be slightly larger – $2 \cdot \log\left(\frac{1}{\varepsilon}\right) + 2$. The algorithm remains unchanged, except for the fact that $\texttt{load}_i(t)$ is now defined as the total processing times of jobs dispatched to machine i till time t . We highlight the main changes in the analysis. Claim 5.1 continues to hold – in inequality (2), the numerator on the right hand side gets replaced by the total processing time of jobs in $J(sT^\star)$. Corollary 5.2 holds with a weaker bound of $(2^{\alpha-1} - 1) \cdot T^\star m_\alpha$ on the number of jobs – this is because the jobs could have size 2. Rest of the arguments remain unchanged. Thus, we get the following result.

Theorem 5.3. *Consider an input instance where all job sizes lie between 1 and 2, and the optimum value is $T^\star$. Then the algorithm described above is $O(\alpha)$-competitive and rejects at most ε-fraction of the jobs.*

5.2 General Processing Times

We now extend the above result to general processing times. Some definitions first. We say that a job j is of class k if $p_j \in [2^k, 2^{k+1})$. We do not know the smallest job size, and so cannot assume (by scaling) that all job sizes are at least 1. Hence, the class of a job could be negative as well. We say that a set of jobs is Δ-*separated* for a positive integer $\Delta > 0$, if there exists an integer k such that the class of any job j in this set belongs to $\{k + i\Delta : i \text{ is an integer}\}$.

Δ-separated jobs: We first assume that the jobs are Δ-separated. Our scheduling algorithm works in two stages. In the first stage, it may violate the objective value by a large amount, but this will get fixed in the second stage. However, we shall ensure that in the first stage, the number of rejected jobs is at most ε-fraction of the total number of jobs. As before, assume that we know the optimal value $T^\star$.

Stage 1: Each machine i maintains a queue of jobs assigned to it. For a class l, let $\texttt{load}_{i,l}(t)$ denote the total processing time of jobs of class l in the queue of machine i at time t. We say that i is *full with respect to class l* at time t if $\texttt{load}_{i,l}(t) \geq \alpha \cdot T^\star$, for some parameter α to be specified later. The dispatch algorithm is as follows: when a job j of class l arrives at time t, dispatch it to the machine $i \in S(j)$ for which $\texttt{load}_{i,l}(t)$ is smallest; unless all such machines are full with respect to class l. If the latter case happens, we reject the job.

Stage 2: This is a pruning step. Note that (and this is important) this step does not affect stage 1 at all. So, while computing $\texttt{load}_{i,l}(t)$ in the algorithm described above, we will assume that no pruning happens. However, without this step, the queue sizes may go much beyond $T^\star$. For a machine i, and time t, let $\texttt{load}_i(t) = \sum_l \texttt{load}_{i,l}(t)$ denote the total load on this machine at time t. In this stage, we do the following for each machine i and time t: if at time t, $\texttt{load}_i(t) > 2\alpha \cdot T^\star$,

we keep removing the largest jobs in the queue of machine i till $\texttt{load}_i(t)$ becomes at most $2\alpha \cdot T^\star$ (we can assume a fixed way of breaking ties).

This completes the description of the algorithm. We now analyze the algorithm. We fix a time $t^\star$ (and we assume that the offline optimum has objective value $T^\star$).

For a fixed class l, the stage 1 algorithm is same as the algorithm for unit size jobs (or when the job sizes differ by a factor of 2 only) described in the previous section. Hence, Theorem 5.3 implies that the total processing time of jobs of class l dispatched to any particular machine is at most $\alpha \cdot T^\star$. Further, the total number of rejected jobs in stage 1 is only ε-times the total number of jobs.

Now, in Stage 2, we ensure that the total load on a machine is at most $2\alpha \cdot T^\star$. We need to bound the number of jobs which get rejected. Fix a machine i. Let the jobs which get dispatched to i (in Stage 1) till $t^\star$ be $j_1, \ldots, j_n$. We assume that the total processing time of these jobs is more than $2\alpha T^\star$, otherwise we do not reject any of these jobs in Stage 2. Also assume that these jobs are arranged in ascending order of processing time. Let $r_{T^\star}$ be the smallest index u for which $p_{j_1} + \cdots + p_{j_u} \geq \alpha \cdot T^\star$ and $r_{2T^\star}$ be the smallest index u such that $p_{j_1} + \cdots + p_{j_u} \geq 2\alpha \cdot T^\star$. Clearly, $r_{2T^\star} \geq r_{T^\star}$. We make some important observations:

(i) During stage 2, we will not remove the jobs $j_1, \ldots, j_{r_{2T^\star}}$. Indeed, in order to remove a job j, there must be at least $2 \cdot \alpha \cdot T^\star$ volume of smaller jobs.

(ii) During stage 2, we will remove all jobs j_u for $u > r_{2T^\star}$. This is true because at time $t^\star$, we have at least $2 \cdot \alpha \cdot T^\star$ volume of jobs which are smaller than this job.

(iii) The class of job $j_{r_{T^\star}}$ is strictly less than that of $j_{r_{2T^\star}}$: Since $j_{r_{T^\star}}$ comes before $j_{r_{2T^\star}}$ in the ordering, the class of $j_{r_{T^\star}}$ is at most that of $j_{r_{2T^\star}}$. Suppose the two jobs are of the same class. Then, the jobs in $j_{r_{T^\star}}, j_{r_{T^\star}+1}, \ldots, j_{r_{2T^\star}}$ are of the same class, and their total processing volume is strictly larger than $\alpha T^\star$. But this is a contradiction – we argued above that for any particular class, we will not dispatch more than $\alpha \cdot T^\star$ volume to a machine. Since the jobs are Δ-separated, class of $j_{r_{T^\star}}$ is at least Δ less than that of $j_{r_{2T^\star}}$.

Let $c_{T^\star}$ and $c_{2T^\star}$ denote the class of the jobs $j_{r_{T^\star}}$ and $j_{r_{2T^\star}}$ respectively. As observed in (iii) above, $c_{2T^\star} \geq c_{T^\star} + \Delta$. We are only rejecting jobs of class $c_{2T^\star}$ or higher (from the queue of machine i). Further, the total number of jobs of a class l that can be rejected is at most $\frac{\alpha \cdot T^\star}{2^l}$, because the total volume of jobs of class l assigned to this machine is at most $\alpha \cdot T^\star$. Thus, the total number of rejected jobs (among those dispatched to machine i) in Stage 2 is at most

$$\sum_{l \geq c_{2T^\star}} \frac{\alpha \cdot T^\star}{2^l} \leq \frac{\alpha \cdot T^\star}{2^{c_{2T^\star}-1}} \leq \frac{\alpha \cdot T^\star}{2^{c_{T^\star}+\Delta-1}} \leq \frac{\varepsilon \cdot \alpha \cdot T^\star}{2^{c_{T^\star}+1}},$$

provided $\Delta = \log\left(\frac{1}{\varepsilon}\right) + 2$. Now, observe that the jobs in $j_1, \ldots, j_{r_{T^\star}}$ have volume at least $\alpha \cdot T^\star$, and their class is at most $c_{T^\star}$. Hence, the number of these jobs, $r_{T^\star}$, is at least

$$\frac{\alpha \cdot T^\star}{2^{c_{T^\star}+1}}.$$

Comparing this with the expression above, we see that at most ε fraction of the jobs are removed in a Stage 2.

Combining the above observations, we get (replacing ε by $\varepsilon/2$ in the argument above)

Lemma 5.4. *Assuming we know the value of the offline optimum and that the jobs are Δ-separated, the above algorithm is $O\left(\log\left(\frac{1}{\varepsilon}\right)\right)$-competitive, where $\Delta = \log\left(\frac{1}{\varepsilon}\right) + 2$. Further, it rejects at most ε-fraction of the jobs.*

Any set of jobs can be partitioned into Δ disjoint sets, each of which is Δ-separated. Running the above algorithm on each such set independently, we get

Corollary 5.5. *Assuming the value of the offline optimum is at most $T^\star$, the algorithm described above assigns load $O\left(\log^2\left(\frac{1}{\varepsilon}\right)\right) \cdot T^\star$ on any machine, and rejects at most ε-fraction of the jobs.*

5.3 Removing the assumption about $T^\star$

So far we have assumed that we know the value of the offline optimum solution. We now show how to relax this assumption. The idea is fairly standard. We start with a small guess for $T^\star$, and double it whenever we realize that the guess was less than the actual offline optimum value. The details are given in Figure 2. For a parameter $T^\star$, let $\mathcal{A}(T^\star)$ denote the algorithm given by Corollary 5.5 where the estimate for the optimal value is $T^\star$. We shall call each iteration of `Step 2` as a *phase*. Note that when we run $\mathcal{A}(T^\star)$ in the beginning of a phase, the algorithm ignores the jobs that have been assigned in the previous phase. It only considers the jobs that arrive next (and hence, quantities like $\texttt{load}_i(t)$ needed by $\mathcal{A}(T^\star)$ do not take into account the jobs which were assigned in previous phases).

Load Balance:

1. Initialize $T^\star \leftarrow p_j$, where j is the first released job.
2. **Repeat**
   (i) Run $\mathcal{A}(T^\star)$ on the jobs which arrive next.
   (ii) If the algorithm rejects more than ε-fraction of the jobs which arrived after the time it started running,
   Stop the algorithm and update $T^\star \leftarrow 2T^\star$.

Figure 2: Algorithm for `LoadBalancing`

Now we analyze the competitive ratio of this algorithm. It is clear that it does not reject more than ε-fraction of the jobs because in each phase, the algorithm rejects at most ε-fraction of the jobs which arrive in that phase. Let $T^{\mathcal{O}}$ denote the objective value of the offline optimum solution. We first argue that the estimate $T^\star$ always stays (almost) below $T^{\mathcal{O}}$.

Claim 5.6. *Except perhaps for the last phase, the value of $T^\star$ is at most $T^{\mathcal{O}}$.*

Proof. Consider the first phase for which $T^\star$ is more than $T^{\mathcal{O}}$. Note that the optimal value for the set of jobs which arrive during this period will be at most $T^{\mathcal{O}}$, which is at most $T^\star$. Corollary 5.5 now implies that $\mathcal{A}(T^\star)$ will not reject more than ε-fraction of the jobs released during this phase, and so this is the last phase of the algorithm. □

The main result now follows easily from the above claim.

Theorem 5.7. *The algorithm* `Load Balance` *is $O\left(\log^2\left(\frac{1}{\varepsilon}\right)\right)$-competitive and rejects at most ε-fraction of the jobs.*

Proof. We have already argued that for any time t, the total number of jobs rejected by the algorithm till time t is at most ε-fraction of the jobs released till this time. The final value of $T^\star$ is at most $2T^{\mathcal{O}}$. Since $T^\star$ increases by a factor of 2 after each phase, Corollary 5.5 and Claim 5.6 imply that the total load assigned to a particular machine is at most

$$O\left(\log^2\left(\frac{1}{\varepsilon}\right)\right)\cdot\left(2T^{\mathcal{O}}+\frac{2T^{\mathcal{O}}}{2}+\frac{2T^{\mathcal{O}}}{4}+\ldots\right)\leq O\left(\log^2\left(\frac{1}{\varepsilon}\right)\right)T^{\mathcal{O}}.$$

This proves the desired result. □

6 Algorithms for the `MaxFlowTime` problem

In this section, we consider the `MaxFlowTime` problem. Again, we shall begin by assuming that we know the optimal value $T^\star$. In Section 6.1, we consider the special case when all jobs have unit size. Here, the algorithm turns out to be a natural one: when a job arrives, send it to the least loaded machine (among the ones it can be processed on). However, the extension to arbitrary job sizes and arbitrary weight turns out to be quite tricky, and although the algorithm remains simple, it does not correspond to a natural idea. We give details of this algorithm in Section 6.2. Finally, we show how to remove the assumption about the knowledge of $T^\star$.

6.1 Unit Job Size

We consider the special case when all jobs have size 1. The algorithm is greedy: for each machine i and time t, it maintains the number of jobs waiting in the queue of machine i at (the beginning of) time t, call this value $\texttt{load}_i(t)$. When a new job j arrives at time t, it gets dispatched to the machine $i \in S_j$ for which $\texttt{load}_i(t)$ is smallest, unless the load on every machine $i \in S_j$ is at least $\alpha T^\star$, where $\alpha = \frac{1}{\varepsilon}$. If the latter happens, we reject the job. Each machines processes jobs in order they are dispatched to it. This completes the description of the algorithm.

Analysis: We first give a way of proving lower bounds on the optimum value. A *machine interval* is defined as a pair (I, i) where I is an interval and i is a machine.

Lemma 6.1. *Let $\mathbf{I}$ be a set of machine intervals and J be a set of jobs. Let α_j be non-negative values assigned to the jobs $j \in J$ such that the following condition is satisfied for all jobs j and machines $i \in S_j$:*

$$\alpha_j \leq |\{(I,i) \in \mathbf{I} : r_j \in I\}|.$$

Then,

$$T^\star \geq \frac{\sum_j \alpha_j - \sum_{(I,i)\in\mathbf{I}} \texttt{length}(I)}{|\mathbf{I}|} \tag{3}$$

Proof. This result follows from weak duality. Consider the following LP relaxation for (the offline optimum of) this problem. For a job j and machine $i \in S_j$, we have a variable x_{ij} which is 1 iff j

is dispatched to i. It is easy to check that the following is a valid LP relaxation:

$$\begin{aligned}
\min \quad & T \\
T &\geq \sum_{j: r_j \in I, i \in S_j} x_{ij} - \texttt{length}(I) \qquad \text{for all machine intervals } (I,i) \\
\sum_{i: i \in S_j} x_{ij} &\geq 1 \qquad \text{for all jobs } j \\
x_{ij} &\geq 0 \qquad \text{for all } i, j
\end{aligned}$$

The dual LP is as follows (variables are α_j and $\beta_{(I,i)}$):

$$\begin{aligned}
\max \qquad & \sum_j \alpha_j - \sum_{(I,i)} \texttt{length}(I) \cdot \beta_{(I,i)} \\
\alpha_j &\leq \sum_{(I,i): r_j \in I, i \in S_j} \beta_{(I,i)} \qquad \text{for all } j \text{ and } i \in S_j \\
\sum_{(I,i)} \beta_{(I,i)} &\leq 1 \\
\alpha_j, \beta_{(I,i)} &\geq 0
\end{aligned}$$

The desired lemma now follows by setting $\beta_{(I,i)} = 1$ for all machine intervals $(I,i) \in \mathbf{I}$. □

We now prove correctness of our algorithm. It is easy to see that the flow-time of any job which does not get rejected is at most $\alpha T^\star$. When a job is dispatched to a machine, the load on the machine is at most $\alpha \cdot T^\star$, and since the machine processes jobs in the order in which they get dispatched to it, this job will finish within $\alpha \cdot T^\star$ time. We now need to bound the number of rejected jobs. Fix a time $t^\star$ – we shall bound the number of rejected jobs till time $t^\star$. Recall that we are assuming that the value of the optimal solution for this input (till time $t^\star$) is at most $T^\star$.

We now give some definitions. For a time t and machine interval (I,i), we shall often abuse notation and say t *lies in* (I,i) (or (I,i) *contains* t) when $t \in I$. Similarly, we shall say that two machine intervals (I_1,i_1) and (I_2,i_2) are *disjoint* if the corresponding intervals I_1 and I_2 are disjoint. Given a set $\mathbf{I}$ of machine intervals, the *set of intervals* in $\mathbf{I}$ refers to the multi-set $\{I : (I,i) \in \mathbf{I}\}$. For each index l, $0 \leq l < \alpha$, and machine i, we define a set (possibly empty) of mutually disjoint machine intervals $\mathcal{I}^{(i,l)}$. The procedure for defining these intervals is described in Figure 3.

It is easy to see that for any value of l and machine i, the set of machine intervals in $\mathcal{I}^{(i,l)}$ are mutually disjoint. We now prove a few simple properties of these sets of intervals. For a job j which gets released at time t and dispatched to machine i, let $\texttt{level}(j)$ denote $\texttt{load}_i(t)$ (without counting the job j itself).

Lemma 6.2. *Fix a machine i and parameter l. The sets of machine intervals $\mathcal{I}^{(i,l)}$ satisfy the following properties:*

- (covering) *Suppose a job j satisfies $\texttt{level}(j) \geq l \cdot T^\star$, and assume $i \in S_j$. Then, there exists a machine interval $(I,i) \in \mathcal{I}^{(i,l)}$ such that $r_j \in I$.*
- (nesting) *Given any machine interval $(I_l,i) \in \mathcal{I}^{(i,l)}$, there exists a machine interval $(I_{l-1},i) \in \mathcal{I}^{(i,l-1)}$ such that I_{l-1} contains I_l.*

Constructing $\mathcal{I}^{(i,l)}$:

1. Initialize $t \leftarrow t^\star$.
2. **Repeat**
   (i) Let t_2 be the highest time before (or equal to) t such that $\texttt{load}_i(t_2) \geq l \cdot T^\star$.
   (ii) Let t_1 be the earliest time before t_2 such that $\texttt{load}_i(t') > (l-1)T^\star$ for all $t' \in (t_1, t_2]$.
   (iii) Add the machine interval $([t_1, t_2], i)$ to $\mathcal{I}^{(i,l)}$.
   (iv) $t \leftarrow t_1$.
3. **Until** No more machine intervals can be added to $\mathcal{I}^{(i,l)}$.

Figure 3: Construction of the set of intervals in $\mathcal{I}^{(i,l)}$

- *For any machine interval $(I_l, i) \in \mathcal{I}^{(i,l)}$, the total number of jobs which get dispatched to i during I_l is at least $\texttt{length}(I_l) + T^\star$.*

Proof. Consider a job j such that $\texttt{level}(j) \geq l \cdot T^\star$, and a machine $i \in S_j$. Then $\texttt{load}_i(r_j) \geq l \cdot T^\star$ as well. Suppose, for the sake of contradiction that there is no machine interval in $\mathcal{I}^{(i,l)}$ such that r_j lies in it. Among the set of intervals in $\mathcal{I}^{(i,l)}$, let I_f be the first interval (if any) to the right of r_j, and suppose s is the starting time of I_f (if I_f does not exist, s is defined as $t^\star$). After adding (I_f, i) to $\mathcal{I}^{(i,l)}$, our algorithm for constructing $\mathcal{I}^{(i,l)}$ would have tried $t_2 = r_j$, and hence there should be a machine interval containing r_j in $\mathcal{I}^{(i,l)}$.

The argument above shows that if for any time t, $\texttt{load}_i(t) \geq l \cdot T^\star$, then there exists a machine interval in $(I, i) \in \mathcal{I}^{(i,l)}$ containing t. Consider $(I_l = (s_l, e_l), i) \in \mathcal{I}^{(i,l)}$. Since $\texttt{load}_i(e_l) \geq lT^\star$, there exists an interval $(I_{l-1} = (s_{l-1}, e_{l-1}), i) \in \mathcal{I}^{(i,l-1)}$ such that $e_l \in I_{l-1}$. Since $\texttt{load}_i(t) > (l-1)T^\star$ for all $t \in (s_l, e_l)$, the left end-point of I_{l-1} will appear before s_l.

The third claim is easy to see. If $I_l = [s, t]$, then $\texttt{load}_i(t) - \texttt{load}_i(s) \geq T^\star$. Further, the machine is busy during $[s, t]$, and so must have processed $t - s$ volume during this period. □

For an index l, let $\mathcal{I}^{(l)}$ denote the union over all machines i of the machine intervals in $\mathcal{I}^{(i,l)}$. Let $L^{(l)}$ denote the total length of the intervals in $\mathcal{I}^{(l)}$, and m_l denote the cardinality of $\mathcal{I}^{(l)}$. Let $J^{(l)}$ be the set of jobs which get dispatched to a machine i during the intervals in $\mathcal{I}^{(i,l)}$.

Let $\mathbf{I}$ denote $\cup_{l=1}^{\alpha} \mathcal{I}^{(l)}$. Let J^R denote the set of rejected jobs. Consider the following assignment of dual variables: if $j \in J^R$, then we set $\alpha_j = \alpha$, otherwise if $\texttt{level}(j) \in [(l-1) \cdot T^\star, lT^\star)$, then we set $\alpha_j = l - 1$. The following claim is easy to see.

Claim 6.3. *The set of machine intervals $\mathbf{I}$ and the α_j values defined above satisfy the conditions of Lemma 6.1.*

Proof. If $j \in J^R$, $\texttt{load}_i(r_j) \geq \alpha \cdot T^\star$ for all $i \in S_j$. Lemma 6.2 now shows that for any $i \in S_j$ and all $l, 1 \leq l \leq \alpha$, there exists a machine interval $(I, i) \in \mathcal{I}^{(i,l)}$ containing r_j. Hence, for any $i \in S_j$, $|\{(I, i) : r_j \in I\}|$ is equal to α. Now suppose $j \notin J^R$. Let k denote α_j. Since $\texttt{level}(j) \geq k \cdot T^\star$, Lemma 6.2 again implies that if $i \in S_j$, then for all l, $1 \leq l \leq k$, there is a machine interval $(I, i) \in \mathcal{I}^{(i,l)}$ such that $r_j \in I$. Hence, the desired result holds in this case as well. □

Using the above claim, we can applying Lemma 6.1 to $\mathbf{I}$ and α_j values to get

$$T^\star \cdot (m_1 + \ldots + m_\alpha) \geq \sum_j \alpha_j - \sum_{l=1}^{\alpha} L^{(l)} \tag{4}$$

Claim 6.4.

$$\sum_{j \notin J^R} \alpha_j \geq \sum_{l>1} L^{(l)} + \sum_{l>1} m_l T^\star.$$

Proof. We first argue that $\sum_{j \notin J^R} \alpha_j \geq \sum_{l>1} |J^{(l)}|$. Consider a job j with $\alpha_j = k$. This implies that $\texttt{level}(j) < (k+1) \cdot T^\star$. So $j \notin J^{(l)}$ for any $l > k+1$. Indeed, suppose j is dispatched to machine i and there is a machine interval $(I, i) \in \mathcal{I}^{(i,l)}$ for some $l > k+1$ such that $r_j \in I$. Then the load on machine i stays above $(k+1) \cdot T^\star$ during I. So $\texttt{level}(j) \geq (k+1) \cdot T^\star$, a contradiction. Therefore, j can only contribute towards $J^{(2)}, \ldots, J^{(k+1)}$ in the right hand side of the above inequality. Hence, $\sum_{j \notin J^R} \alpha_j \geq \sum_{l>1} |J^{(l)}|$. The claim now follows from Lemma 6.2 (third part). □

We can now show that the number of rejected jobs is small.

Corollary 6.5. *The total number of rejected jobs is at most ε times the total number of jobs. Further, the flow-time of any job is at most $T^\star/\varepsilon$.*

Proof. Applying Claim 6.4 to inequality (4), we get

$$T^\star \cdot m_1 + L^{(1)} \geq \sum_{j \in J^R} \alpha_j = \alpha \cdot |J^R|.$$

Observe that the number of jobs is at least $T^\star \cdot m_1 + L^{(1)}$ – this follows from the third part of Lemma 6.2 and the fact that for any machine i, all intervals in $\mathcal{I}^{(i,1)}$ are disjoint. If we pick $\alpha = \frac{1}{\varepsilon}$, the desired result follows. The second statement follows from the fact that the queue size never exceeds $\alpha T^\star$. □

6.2 Extension to `WtdMaxFlowTime`

We now extend the above result to the case of arbitrary job sizes and associated weights. We can assume without loss of generality that weight of a job is a power of 2 – this will affect the objective function by a factor of 2 only. Further, the weight of rejected jobs could be off by a factor of 2 – but we can replace ε by $\varepsilon/2$ to take care of this (this will affect the flow-time of a job only by a constant factor).

We shall describe the scheduling algorithm in two stages. First we describe an algorithm $\mathcal{A}$ – this algorithm may not ensure that flow-time of all jobs (which are not rejected) are small, but it will ensure that the queue sizes at all times (on any machine) will be small. Our actual algorithm, which we call $\mathcal{B}$, will use $\mathcal{A}$ in the *background*. It will emulate $\mathcal{A}$, but will periodically reject more jobs, and may prefer to process higher density jobs at a time slot (as compared to $\mathcal{A}$).

Algorithm $\mathcal{A}$: Before we describe the algorithm $\mathcal{A}$, we give some definitions. These definitions group jobs into various *classes* depending on their processing requirement or weight or density. We say that a job j is of *size class* p if $p_j \in [2^p, 2^{p+1})$. We say that it is of *weight class* w if $w_j = 2^w$ (recall that weight of a job is a power of 2), and of density class d if its density, i.e., w_j/p_j lies in the range $[2^d, 2^{d+1})$. Observe that if the size, weight and density classes of a job are p, w, and d respectively, then $d = w - p - 1$.

Note that we cannot assume any lower bound on p_j or w_j, and so, the classes could be negative. Define $\texttt{type}(j) = (w, d)$, if its density class is d, and its weight class is w.

Each machine i maintains a queue for jobs of a particular type. For a machine i, time t, and pair (w, d), let $Q_{i,w,d}(t)$ denote the jobs of type (w, d) waiting in the queue of machine i at time t.

Let $\texttt{load}_{i,w,d}(t)$ denote the total weighted remaining processing time of the jobs in $Q_{i,w,d}(t)$ – if a job j has remaining processing time p'_j, then its *weighted remaining processing time* is defined as $w_j p'_j$.

When a job j of type (w,d) arrives at time t, we dispatch it to the machine $i \in S_j$ for which $\texttt{load}_{i,w,d}(t)$ is minimum, unless for all $i \in S_j$, $\texttt{load}_{i,w,d}(t) + p_j.2^w \geq \alpha^2 \cdot T^\star$. If the latter case happens, we reject this job. Here $\alpha = \frac{76}{\varepsilon}$. It remains to specify which job is processed at any time by a machine.

For a time t and machine i, let $(w_i^\star(t), d_i^\star(t))$ be the pair (w,d) with the highest $2^d \cdot \texttt{load}_{i,w,d}(t)$ value. We process the earliest released job from the queue $Q_{i,w_i^\star(t),d_i^\star(t)}(t)$ on machine i at time t. Assume a fixed rule of breaking ties.

This completes the description of the algorithm $\mathcal{A}$. We note a few important aspects: (i) For a machine i and pair (w,d), the algorithm always prefers the earliest released job in the queue of type (w,d) jobs. Hence, at any time t, there will be at most one job in $Q_{i,w,d}(t)$ which is partially processed, (ii) The policy which decides which job to process at a time t on a machine i balances two aspects: it prefers jobs of higher density, but also prefers jobs for which the corresponding queue is close to the maximum limit – the total weighted remaining processing times of jobs in the queue $Q_{i,w,d}(t)$ should not exceed a constant times $T^\star$.

Algorithm $\mathcal{B}$: Now we describe the actual scheduling algorithm $\mathcal{B}$. When a job j arrives at time t, it is dispatched according to $\mathcal{A}$: if $\mathcal{A}$ rejects this job, $\mathcal{B}$ also rejects it; and if $\mathcal{A}$ dispatched it to machine i, then $\mathcal{B}$ also dispatches this job to i. Now, we describe the processing policy for a fixed machine i. Consider a time t. Let $d^{\mathcal{A}}(t)$ denote the density class of the job processed by $\mathcal{A}$ at time t. Then, $\mathcal{B}$ processes the following job at time t: if there is a job of density class at least $d^{\mathcal{A}}(t) + \log\left(\frac{1}{\varepsilon}\right)$ in the queue of machine i at time t, then $\mathcal{B}$ processes any such job; otherwise it processes the job of density class $d^{\mathcal{A}}(t)$ with the highest weight class (it prefers the earliest released job in case of ties). Also note that if the second case happens and there is no job of density $d^{\mathcal{A}}(t)$ in the queue of machine i at time t (in $\mathcal{B}$), we can process any job at this time.

The algorithm $\mathcal{B}$ may reject some more jobs. For a weight class w, we divide the time line into segments of length $\frac{\alpha^2 T^\star}{\varepsilon 2^w}$. Suppose a job of weight class w gets released during such a segment S, and let S' be the segment immediately to the right of S. If the job does not complete processing by the end of S', the algorithm $\mathcal{B}$ rejects the job. This completes the description of $\mathcal{B}$.

6.2.1 Analysis

We now analyze the scheduling algorithm. We first consider the algorithm $\mathcal{A}$ and prove that the queues remain bounded in size.

Analysis for algorithm $\mathcal{A}$: As in the case of unit sized jobs, we first give a lower bound for $T^\star$. A *weighted machine interval* is defined as a triplet (I,i,w) where I is a time interval, i is a machine, and w denotes a weight class. We say that a time t lies in (or belongs to) a weighted machine interval (I,i,w) if $t \in I$. Similarly, we say that two machine intervals are disjoint (or nested) if the same holds for the associated time intervals.

Lemma 6.6. *Let $\mathbf{I}$ be a set of weighted machine intervals. Let α_j be non-negative values assigned to the jobs $j \in J$ such that the following condition is satisfied for all jobs j and machines $i \in S_j$:*

$$\frac{\alpha_j}{p_j} \leq |\{(I,i,w') \in \mathbf{I} : r_j \in I, w' \leq w\}|, \tag{5}$$

where w denotes the weight class of j, i.e., $w_j = 2^w$. Then,

$$T^\star \geq \frac{\sum_j \alpha_j - \sum_{(I,i,w)\in \mathbf{I}} \texttt{length}(I)}{\sum_{(I,i,w)\in \mathbf{I}} \frac{1}{2^w}} \tag{6}$$

Proof. The proof again follows from LP duality. It is easy to check that the following is a valid LP relaxation – here x_{ij} is a variable which is 1 if job j is assigned to machine i, 0 otherwise.

$$\begin{aligned}
&\min \quad T \\
\sum_{j:r_j\in I, i\in S_j, w_j\geq 2^w} p_j \cdot x_{ij} - \texttt{length}(I) &\leq \frac{T}{2^w} \qquad \text{for all weighted machine intervals } (I,i,w) \\
\sum_{i:i\in S_j} x_{ij} &\geq 1 \qquad \text{for all jobs } j \\
x_{ij} &\geq 0 \qquad \text{for all jobs } j \text{ and machines } i,\ i\in S_j
\end{aligned}$$

The dual LP is as follows (variables are α_j and $\beta_{(I,i,w)}$):

$$\begin{aligned}
\max \qquad & \sum_j \alpha_j - \sum_{(I,i,w)} \texttt{length}(I)\cdot \beta_{(I,i,w)} \\
\frac{\alpha_j}{p_j} &\leq \sum_{(I,i,w):r_j\in I, 2^w\leq w_j} \beta_{(I,i,w)} \qquad \text{for all jobs } j \text{ and machines } i\in S_j \\
\sum_{(I,i,w)} \frac{\beta_{(I,i,w)}}{2^w} &\leq 1 \\
\alpha_j, \beta_{(I,i,w)} &\geq 0
\end{aligned}$$

The lemma now follows by setting $\beta_{(I,i,w)}$ to 1 if $(I,i,w)\in\mathbf{I}$, 0 otherwise. □

Now we bound the weight of jobs rejected by $\mathcal{A}$. We fix a time $t^\star$ – this is the time by which all the jobs finish processing. For rest of the discussion, we shall also fix a density class $d^\star$. We shall first bound the total weight of jobs of density class $d^\star$ which are rejected by the algorithm $\mathcal{A}$. Finally, we shall take a sum over all values of $d^\star$ to bound the total weight of jobs rejected by $\mathcal{A}$.

We begin with some definitions. Define $\Delta = \log(\alpha)$ (assume wlog that α is a power of 2). We shall be interested in jobs whose density class lies in the range $[d^\star, d^\star+\Delta]$. For an integer l, let γ_l denote $4l\alpha T^\star$. We will ensure that $\alpha/8 \leq l \leq \alpha/4$, and hence, γ_l will lie in the range $[\alpha^2 T^\star/2, \alpha^2 T^\star]$. For a machine i, density class d and time t, let $\texttt{load}_{i,d}(t)$ denote the maximum, over all pairs (w,d) of the total weighted remaining processing time of jobs in the queue $Q_{i,w,d}$, i.e.,

$$\texttt{load}_{i,d}(t) = \max_w \texttt{load}_{i,w,d}(t).$$

Consider a job j of type (w,d) which gets dispatched to machine i. We define $\texttt{level}(j)$ as $\texttt{load}_{i,w,d}(r_j) + w_j \cdot p_j$ – this is the load it sees on machine i (including its own weighted processing size). For a parameter l, we define a set of jobs $J^{(l,d^\star)}$ as follows: a job j of density class $d \in [d^\star, d^\star+\Delta]$ lies belongs to $J^{(l,d^\star)}$ iff $\texttt{level}(j) \geq \frac{\gamma_l}{2^{d-d^\star}}$.

For each machine i and parameter l lying the range as mentioned above, we define a set of disjoint weighted machine intervals $\mathcal{I}^{(i,l,d^\star)}$. The procedure for this is given in Figure 4. For a machine i and time t, $\texttt{maxload}_{i,d^\star}(t)$ is defined as $\max_{d\in[d^\star,d^\star+\Delta]} 2^{d-d^\star}\cdot \texttt{load}_{i,d}(t)$.

We now prove the analogue of Lemma 6.2.

Constructing $\mathcal{I}^{(i,l,d^\star)}$:

1. Initialize t as $t^\star$ and $\mathcal{I}^{(i,l,d^\star)} \leftarrow \emptyset$
2. **Repeat**
   (i) Let t_2 be the highest time before (or equal to) t such that $\texttt{maxload}_{i,d^\star}(t_2) \geq \gamma_l$.
   (ii) Let t_1 be the earliest time such that $\texttt{maxload}_{i,d^\star}(t') > \gamma_{l-1}$ for all $t' \in (t_1, t_2]$.
   (iii) Let w be the smallest weight class of a job $j \in J^{(l-1,d^\star)}$ such that $r_j \in (t_1, t_2]$ and $i \in S_j$.
   (iv) Add the weighted machine interval $([t_1, t_2], i, w)$ to $\mathcal{I}^{(i,l,d^\star)}$.
   (v) $t \leftarrow t_1$.
   Until No more intervals can be added to $\mathcal{I}^{(i,l,d^\star)}$.

Figure 4: Construction of the set of intervals in $\mathcal{I}^{(i,l,d^\star)}$ for a machine i and parameter l.

Lemma 6.7. *The sets of intervals in $\mathcal{I}^{(i,l,d^\star)}$ satisfy the following properties:*

- (covering) *Consider a job $j \in J^{(l,d^\star)}$ and a machine $i \in S_j$ Then there exists a weighted machine interval $(I, i, w) \in \mathcal{I}^{(i,l-1,d^\star)}$ such that $r_j \in I$ and $w_j \geq 2^w$.*
- (nesting) *Given any weighted machine interval $(I_l, i, w) \in \mathcal{I}^{(i,l,d^\star)}$, there exists a weighted machine interval $(I_{l-1}, i, w') \in \mathcal{I}^{(i,l-1,d^\star)}$ for some $w' \leq w$ such that I_{l-1} contains I_l.*
- (processing) *For any weighted machine interval $(I, i, w) \in \mathcal{I}^{(i,l,d^\star)}$, the total duration in I during which jobs of density class in $[d^\star, d^\star + \Delta]$ are processed on machine i is at most the total processing size of jobs in $J^{(l-1,d^\star)} \cap \{j : w_j \geq 2^w, r_j \in I\}$ which get dispatched to i during I.*

Proof. Consider a job $j \in J^{(l,d^\star)}$ of type (w, d), and a machine $i \in S_j$:

$$\texttt{load}_{i,d}(r_j) \geq \texttt{load}_{i,w,d}(r_j) \geq \texttt{level}(j) - w_j \cdot p_j \geq \frac{\gamma_l}{2^{d-d^\star}} - 4T^\star \geq \frac{\gamma_l}{2^{d-d^\star}} - \frac{4\alpha T^\star}{2^{d-d^\star}} = \frac{\gamma_{l-1}}{2^{d-d^\star}},$$

where the second inequality follows by our dispatch rule and the third inequality follows from the fact that $T^\star \geq w_j \cdot p_j$ for any j. Therefore, $\texttt{maxload}_{i,d^\star}(r_j) \geq \gamma_{l-1}$. Suppose, for the sake of contradiction that there is no weighted machine interval in $\mathcal{I}^{(i,l-1,d^\star)}$ containing r_j. Let I_f be the first interval (if any) in $\mathcal{I}^{(i,l-1,d^\star)}$ to the right of r_j, and suppose s is the starting time of I_f (if I_f does not exist, s is defined as $t^\star$). After adding I_f to $\mathcal{I}^{(i,l-1,d^\star)}$, our algorithm for constructing $\mathcal{I}^{(i,l-1,d)}$ would have tried $t_2 = r_j$, and hence there should be a weighted machine interval in $\mathcal{I}^{(i,l-1,d^\star)}$ containing r_j. This is a contradiction. Therefore, there is a $(I, i, w') \in \mathcal{I}^{(i,l-1,d^\star)}$ containing r_j. Since $j \in J^{(l-1,d^\star)}$, and $r_j \in I, i \in S_j$, it must be the case that $w' \leq w$. This proves the first part of the lemma.

The argument above shows that if for any time t, machine i and density class $d \in [d^\star, d^\star + \Delta]$, $\texttt{load}_{i,d}(t) \geq \frac{\gamma_l}{2^{d-d^\star}}$, then there exists a weighted machine interval $(I, i, w) \in \mathcal{I}^{(i,l-1,d^\star)}$ containing t. Suppose $(I_l = (s_l, e_l), i, w) \in \mathcal{I}^{(i,l,d^\star)}$. Since $\texttt{maxload}_{i,d^\star}(e_l) \geq \gamma_l$, there is a density class $d \in [d^\star, d^\star + \Delta]$ such that $\texttt{load}_{i,d}(e_l) \geq \frac{\gamma_l}{2^{d-d^\star}}$. Hence, there exists $(I_{l-1} = (s_{l-1}, e_{l-1}), i, w') \in \mathcal{I}^{(i,l-1,d^\star)}$ containing e_l. Since $\texttt{maxload}_{i,d^\star}(t) > \gamma_{l-1}$ for all $t \in (s_l, e_l)$, the left end-point of I_{l-1} will appear before s_l. Moreover, $w' \leq w$ (follows from the definition of w or w' in Step 2(iii) of Figure 4 and the fact that I_l is contained in I_{l-1}). This proves the second part of the lemma.

It remains to prove the third part. Consider a weighted machine interval $(I, i, w) \in \mathcal{I}^{(i,l,d^\star)}$.

Claim 6.8. *Suppose we process a job j of type (w,d) on machine i at a particular time $t \in I$. Then it must be the case that $d \geq d^\star$ and*

$$\texttt{load}_{i,d}(t) = \texttt{load}_{i,w,d}(t) > \frac{\gamma_{l-1}}{2^{d-d^\star}}.$$

Proof. Suppose $d < d^\star$. We know that there is a density class $d' \in [d^\star, d^\star+\Delta]$ such that $\texttt{load}_{i,d'}(t) > \frac{\gamma_{l-1}}{2^{d'-d^\star}}$. The job processing rule for $\mathcal{A}$ implies that $2^d \cdot \texttt{load}_{i,w,d}(t) \geq 2^{d'} \cdot \texttt{load}_{i,d'}(t) > 2^{d^\star}\gamma_{l-1}$, and so, $\texttt{load}_{i,w,d}(t) > 2\gamma_{l-1} \geq \alpha^2 T^\star$, because $\gamma_{l-1} \geq \alpha T^\star/2$. But this is a contradiction (because of the job dispatch policy of $\mathcal{A}$). Hence, $d \geq d^\star$.

If $\texttt{load}_{i,d}(t) > \texttt{load}_{i,w,d}(t)$, then let w' be the weight class for which $\texttt{load}_{i,d}(t) = \texttt{load}_{i,w',d}(t)$. But then $2^d \cdot \texttt{load}_{i,w,d}(t) < 2^d \cdot \texttt{load}_{i,w',d}(t)$, and so $\mathcal{A}$ cannot process j on machine i at time t. Similarly, if d' is as above, then it must be the case that

$$2^d \cdot \texttt{load}_{i,w,d}(t) \geq 2^{d'} \cdot \texttt{load}_{i,d'}(t) > \gamma_{l-1} \cdot 2^{d^\star}.$$

This implies the claim. □

Now consider a density class $d \in [d^\star, d^\star+\Delta]$, and a weight class w'. Let t_f be the first time in I at which $\mathcal{A}$ processes a job j of type (w',d). The claim above shows that $\texttt{load}_{i,d}(t_f) = \texttt{load}_{i,w',d}(t_f) > \frac{\gamma_{l-1}}{2^{d-d^\star}}$. Further, $\texttt{load}_{i,w',d}(t)$ remains at least $\frac{\gamma_{l-1}}{2^{d-d^\star}}$ after t_f till the end of the interval I – indeed, Claim 6.8 says that if $\texttt{load}_{i,w',d}(t)$ is at most $\frac{\gamma_{l-1}}{2^{d-d^\star}}$ for some time $t \in I$, then the algorithm does not process a job of type (w',d) at this time on machine i. Hence, $\texttt{load}_{i,w',d}(t)$ will not go below $\frac{\gamma_{l-1}}{2^{d-d^\star}}$ after t_f (during I). Also, for any time t during $[t_1, t_f)$, $\texttt{load}_{i,w',d}(t) < \frac{\gamma_{l-1}}{2^{d-d^\star}}$, where t_1 is the left end-point of I (by definition of t_f and the fact that this statement holds for t_1). Let V denote the total volume of time during I when we process a job of type (w',d). Then it must happen that at least V volume of jobs of type (w',d) are released during $[t_f, t_e]$, where t_e is the right end-point of I – if this does not happen then we will end up processing a job of type (w',d) at a time t in I even when $\texttt{load}_{i,w',d}(t) < \frac{\gamma_{l-1}}{2^{d-d^\star}}$, a contradiction. For any job j' of type (w',d) released during $[t_f, t_e]$, $\texttt{load}_{i,w',d}(r_j) \geq \frac{\gamma_{l-1}}{2^{d-d^\star}}$, and so, $j' \in J^{(l-1,d^\star)}$. Recall that the weighted machine interval was denoted by (I,i,w) – by definition of w, it must be satisfy $w \leq w'$. Summing over all w' gives us the lemma. □

We now define a set of weighted machine intervals $\mathbf{I}$ and dual values α_j for all jobs j which will satisfy the conditions of Lemma 6.6. Let $\mathbf{I}$ be the set of weighted machine intervals defined above, i.e., $\cup_{l=\alpha/8}^{\alpha/4} \cup_i \mathcal{I}^{(i,l,d^\star)}$ – note that this is a multi-set, i.e., if a weighted machine interval (I,i,w) appears in several of the sets $\mathcal{I}^{(i,l,d^\star)}$, it is counted these many times. For each job $j \in \cup_{l=\alpha/8}^{\alpha/4} J^{(l,d^\star)}$ which does not get rejected, define $\alpha_j = p_j \cdot (l_j - \alpha/8)$, where l_j is the largest value l such that $j \in J^{(l,d^\star)}$. Further, if j is a job of density class $d^\star$ gets rejected, we set $\alpha_j = p_j \cdot \alpha/8$. For remaining jobs, we set $\alpha_j = 0$.

Claim 6.9. *The set of intervals $\mathbf{I}$ and the values α_j defined above satisfy the feasibility conditions* (5).

Proof. Consider a job j for which $\alpha_j = p_j \cdot (l_j - \alpha/8)$, and a machine $i \in S_j$. Lemma 6.7 shows that there are weighted machine intervals $(I_l, i, w_l) \in \mathcal{I}^{(i,l,d^\star)}$ for $l = \alpha/8, \ldots, l_j - 1$, such that $r_j \in I_l$ and $2^{w_l} \leq w_j$. Thus, the conditions (5) are satisfied for j. Similarly, if j is a job of density class $d^\star$ which gets rejected, and $i \in S_j$, then $\texttt{level}(j) \geq \alpha^2 T^\star \geq \gamma_l$, for $l = \alpha/4$. Thus, $j \in J^{(\alpha/4,d^\star)}$, and so, the same argument as above applies here as well. □

Claim 6.9 implies that we can apply Lemma 6.6. We give some notation first. For a parameter l and weight class w, let $\mathbf{I}^{(l,w,d^\star)}$ denote the following set of weighted machine intervals:

$$\{(I,i,w) : (I,i,w) \in \cup_{i'} \mathcal{I}^{(i',l,d^\star)}\}.$$

Let $m_{l,w,d^\star}$ denote $|\mathbf{I}^{(l,w,d^\star)}|$, and $L^{(l,w,d^\star)}$ denote the total length of (associated intervals in) the weighted machine intervals in $\mathbf{I}^{(l,w,d^\star)}$. Applying Lemma 6.6, we get

$$T^\star \cdot \sum_w \sum_{l=\alpha/8}^{\alpha/4} \frac{m_{l,w,d^\star}}{2^w} \geq \sum_j \alpha_j - \sum_w \sum_{l=\alpha/8}^{\alpha/4} L^{(l,w,d^\star)}. \tag{7}$$

Claim 6.10. *Let $J^R(d^\star)$ denote the jobs of density class $d^\star$ which get rejected. Also let $P_{\geq d^\star+\Delta}$ denote the total processing time of jobs of density class higher than $d^\star+\Delta$. Then,*

$$\sum_{j \notin J^R(d^\star)} \alpha_j \geq \sum_w \sum_{l=\alpha/8+1}^{\alpha/4} L^{(l,w,d^\star)} - \alpha/8 \cdot P_{\geq d^\star+\Delta}.$$

Proof. We split $L^{(l,w,d^\star)}$ into two parts – let $L^{(l,w,d^\star)'}$ denote the volume during intervals in $\mathbf{I}^{(l,w,d^\star)}$ where $\mathcal{A}$ processes a job of density class lying in the set $[d^\star, d^\star+\Delta]$, and $L^{(l,w,d^\star)''}$ be the volume during intervals in $\mathbf{I}^{(l,w,d^\star)}$ where $\mathcal{A}$ processes a job of class density higher than $d^\star+\Delta$. Claim 6.8 implies that $L^{(l,w,d^\star)} = L^{(l,w,d^\star)'} + L^{(l,w,d^\star)''}$.

Lemma 6.7 (third part) implies that $\sum_w L^{(l,w,d^\star)'} \leq \sum_{j \in J^{(l-1,d^\star)}, j \notin J^R(d^\star)} p_j$ (the intervals in $\mathbf{I}^{(w,l,d^\star)} \cap \mathcal{I}^{(i,l,d^\star)}$ are disjoint). Summing over $l = \alpha/8+1, \ldots, \alpha/4$, we get

$$\sum_{l=\alpha/8+1}^{\alpha/4} \sum_w L^{(l,w,d^\star)'} \leq \sum_{l=\alpha/8+1}^{\alpha/4} \sum_{j \in J^{(l-1,d^\star)}, j \notin J^R(d^\star)} p_j = \sum_{j \notin J^R(d^\star)} p_j \cdot (l_j - \alpha/8) = \sum_{j \notin J^R(d^\star)} \alpha_j$$

Now we bound $L^{(l,w,d^\star)''}$. For a fixed l, $\sum_w L^{(l,w,d^\star)''}$ is at most $P_{\geq d^\star+\Delta}$. Thus,

$$\sum_{l=\alpha/8+1}^{\alpha/4} \sum_w L^{(l,w,d^\star)''} \leq \alpha/8 \cdot P_{\geq d^\star+\Delta}.$$

Combining the above inequalities gives us the desired result. □

Hence, inequality (7) can be simplified as

$$T^\star \cdot \sum_w \sum_{l=\alpha/8}^{\alpha/4} \frac{m_{l,w,d^\star}}{2^w} + \alpha/8 \cdot P_{\geq d^\star+\Delta} + \sum_w L^{(\alpha/8,w,d^\star)} \geq \sum_{j \in J^R(d^\star)} \alpha_j = \sum_{j \in J^R(d^\star)} \frac{\alpha \cdot p_j}{8} \tag{8}$$

We simplify the above expression to bound the weight of jobs of density class $d^\star$ which get rejected. Some more notation. For a weight class w and density class d, let $V_{w,d}$ denote the total weighted processing size of jobs of type (w,d), and let $W_{w,d}$ be the total weight of such jobs, i.e.,

$$V_{w,d} = \sum_{j:\texttt{type}(j)=(w,d)} w_j p_j, \quad W_{w,d} = \sum_{j:\texttt{type}(j)=(w,d)} w_j.$$

For parameters $a, b, a \leq b$, let $V_{w,d}(a,b)$ and $W_{w,d}(a,b)$ denote the corresponding sum for jobs j with $\texttt{level}(j) \in [a,b]$, i.e,

$$V_{w,d}(a,b) = \sum_{j:\texttt{type}(j)=(w,d),\texttt{level}(j)\in[a,b]} w_j p_j, \quad W_{w,d}(a,b) = \sum_{j:\texttt{type}(j)=(w,d),\texttt{level}(j)\in[a,b]} w_j.$$

Claim 6.11.

$$T^\star \cdot \sum_w \sum_{l=\alpha/8}^{\alpha/4} \frac{m_{l,w,d^\star}}{2^w} + \sum_w L^{(\alpha/8,w,d^\star)} \leq \sum_{d\in[d^\star,d^\star+\Delta]} \sum_w \frac{1}{2\alpha 2^{d^\star}} W_{w,d} + P_{\geq d^\star},$$

where $P_{\geq d^\star}$ denote the total processing time of jobs of density class at least $d^\star$.

Proof. First we consider $\sum_w L^{(\alpha/8,w,d^\star)}$ – this is just the total length of the weighted machine intervals $\cup_i \mathcal{I}^{(i,\alpha/8,d^\star)}$. Note that these weighted machine intervals are disjoint and Lemma 6.7 implies that we will only process jobs of density class at least $d^\star$ during these intervals. Hence, $\sum_{L^{(\alpha/8,w,d^\star)}}$ is at most $P_{\geq d^\star}$.

Finally, we consider the term $T^\star \cdot \sum_w \sum_{l=\alpha/8}^{\alpha/4} \frac{m_{l,w,d^\star}}{2^w}$. We shall upper bound each term of the summation as follows. Consider a machine interval $(I,i,w) \in \mathcal{I}^{(i,l,d^\star)}$. Let $I = [s,t]$. We know that there exits a pair (w,d) such that $\texttt{load}_{i,w,d}(s) < \frac{\gamma_{l-1}}{2^{d-d^\star}}$, but $\texttt{load}_{i,w,d}(t) \geq \frac{\gamma_l}{2^{d-d^\star}}$. Let j_1 be the first job of type (w,d) dispatched to i during I for which $\texttt{level}(j_1) \geq \frac{\gamma_{l-1}}{2^{d-d^\star}}$ and j_2 be the first such job for which $\texttt{level}(j_2) \geq \frac{\gamma_l}{2^{d-d^\star}}$. All jobs of type (w,d) dispatched to i after j_1 and before j_2 will have level in the range $\left(\frac{\gamma_{l-1}}{2^{d-d^\star}}, \frac{\gamma_l}{2^{d-d^\star}}\right)$, because as argued in the proof of Lemma 6.7(third part), once $\texttt{load}_{i,w,d}(t)$ goes above $\frac{\gamma_l}{2^{d-d^\star}}$, it will never go below this quantity till the end of I. Further the total weighted volume of such jobs will be at least

$$\frac{\gamma_l}{2^{d-d^\star}} - \frac{\gamma_{l-1}}{2^{d-d^\star}} - w_{j_1}p_{j_1} - w_{j_2}p_{j_2} \geq \frac{4\alpha T^\star}{2^{d-d^\star}} - 2T^\star \geq \frac{2\alpha T^\star}{2^{d-d^\star}},$$

where the last inequality follows from the fact that $\alpha \geq 2^{d-d^\star}$. Thus, we get

$$\frac{T^\star}{2^w} \leq \frac{2^{d-d^\star}}{2\alpha 2^w} \cdot \sum_{j:r_j\in I,\texttt{type}(j)=(w,d),\texttt{level}(j)\in\left(\frac{\gamma_{l-1}}{2^{d-d^\star}},\frac{\gamma_l}{2^{d-d^\star}}\right)} w_j p_j$$

Summing over all weighted machine intervals in $\cup_i \mathcal{I}^{(i,l,d^\star)}$, we get

$$\sum_w \frac{T^\star \cdot m_{l,w,d^\star}}{2^w} \leq \sum_{d\in[d^\star,d^\star+\Delta]} \sum_w \frac{2^{d-d^\star}}{2\alpha 2^w} \cdot V_{w,d}\left(\frac{\gamma_{l-1}}{2^{d-d^\star}}, \frac{\gamma_l}{2^{d-d^\star}}\right).$$

Summing over all l and using the fact that if $a \leq b \leq c$, then $V(a,b) + V(b,c) = V(a,c)$, we get

$$\begin{aligned}\sum_{l=\alpha/8}^{\alpha/4} \sum_w \frac{T^\star \cdot m_{l,w,d^\star}}{2^w} &\leq \sum_{d\in[d^\star,d^\star+\Delta]} \sum_w \frac{2^{d-d^\star}}{2\alpha 2^w} \cdot V_{w,d}\left(\frac{\gamma_{\alpha/8-1}}{2^{d-d^\star}}, \frac{\gamma_{\alpha/4}}{2^{d-d^\star}}\right) \\ &\leq \sum_{d\in[d^\star,d^\star+\Delta]} \sum_w \frac{2^{d-d^\star}}{2\alpha 2^w} \cdot V_{w,d} \leq \sum_{d\in[d^\star,d^\star+\Delta]} \frac{1}{2\alpha 2^{d^\star}} \cdot \sum_w W_{w,d},\end{aligned}$$

where the last inequality follows from the fact that for any job j of type (w,d) $\frac{2^w}{2^{d+1}} < p_j \leq \frac{2^w}{2^d}$ and so,

$$\frac{2^d}{2^w} V_{w,d} \leq \sum_{j:\texttt{type}(j)=(w,d)} w_j/p_j \cdot p_j = W_{w,d}.$$

Thus we have shown the desired result. □

We are now ready to bound the weight of jobs rejected by $\mathcal{A}$.

Lemma 6.12. *The total weight of jobs rejected by $\mathcal{A}$ is at most ε times the total weight of all the jobs, provided we pick $\alpha = \frac{76}{\varepsilon}$.*

Proof. We first bound the total weight of jobs of density class $d^\star$ which get rejected by $\mathcal{A}$. We have

$$\begin{aligned}
\sum_{j \in J^R(d^\star)} w_j \leq 2^{d^\star+1} \cdot \sum_{j \in J^R(d^\star)} p_j &\overset{(8)}{\leq} \frac{16 \cdot 2^{d^\star} T^\star}{\alpha} \cdot \sum_w \sum_{l=\alpha/8}^{\alpha/4} \frac{m_{l,w,d^\star}}{2^w} + 2^{d^\star+1} \cdot P_{\geq d^\star+\Delta} + \frac{16 \cdot 2^{d^\star}}{\alpha} \sum_w L^{(\alpha/8,w,d^\star)} \\
&\overset{Claim\ 6.11}{\leq} \sum_{d \in [d^\star, d^\star+\Delta]} \sum_w \frac{8}{\alpha^2} W_{w,d} + \frac{16 \cdot 2^{d^\star}}{\alpha} P_{\geq d^\star} + 2^{d^\star+1} \cdot P_{\geq d^\star+\Delta} \\
&= \frac{8}{\alpha^2} \cdot \sum_{d=d^\star}^{d^\star+\Delta} \sum_{j \in J_d} w_j + \frac{16 \cdot 2^{d^\star}}{\alpha} \sum_{d \geq d^\star} \sum_{j \in J_d} p_j + 2^{d^\star+1} \sum_{d \geq d^\star+\Delta} \sum_{j \in J_d} p_j
\end{aligned}$$

where J_d denotes the jobs of density class d. Summing over all values of $d^\star$, the total weight of rejected jobs can be expressed as

$$\begin{aligned}
\sum_{d^\star} \sum_{j \in J^R(d^\star)} w_j &\leq \frac{8}{\alpha^2} \sum_{d^\star} \sum_{d=d^\star}^{d^\star+\Delta} \sum_{j \in J_d} w_j + \sum_{d^\star} \frac{16 \cdot 2^{d^\star}}{\alpha} \sum_{d \geq d^\star} \sum_{j \in J_d} p_j + \sum_{d^\star} 2^{d^\star+1} \sum_{d \geq d^\star+\Delta} \sum_{j \in J_d} p_j \\
&\leq \frac{8\Delta}{\alpha^2} \sum_d \sum_{j \in J_d} w_j + \sum_d \sum_{j \in J_d} p_j \sum_{d^\star \leq d} \frac{16 \cdot 2^{d^\star}}{\alpha} + \sum_d \sum_{j \in J_d} p_j \sum_{d^\star \leq d-\Delta} 2^{d^\star+1} \\
&\leq \frac{8\Delta}{\alpha^2} \sum_j w_j + \frac{64}{\alpha} \sum_d \sum_{j \in J_d} p_j 2^{d+1} + \frac{4}{\alpha} \sum_d \sum_{j \in J_d} p_j 2^d \\
&\leq \left(\frac{8\Delta}{\alpha^2} + \frac{64}{\alpha} + \frac{4}{\alpha} \right) \sum_j w_j \ \leq \ \varepsilon \cdot \sum_j w_j,
\end{aligned}$$

if we pick $\alpha = \frac{76}{\varepsilon}$. □

Thus we have shown the main theorem of this section.

Theorem 6.13. *The algorithm $\mathcal{A}$ rejects jobs of total weight at most ε times the total weight of all jobs, and ensures that for any machine i, time t, and pair (w,d), the total weighted remaining processing time of jobs of type (w,d) at time t on machine i is at most $\alpha^2 T^\star$. Further, $\mathcal{A}$ is an immediate dispatch algorithm which rejects jobs on arrival only.*

Analysis for algorithm $\mathcal{B}$: Now we analyze the algorithm $\mathcal{B}$. It is clear that the weighted flow-time of any job j is at most $2\beta T^\star$, where β denotes $\frac{\alpha^2}{\varepsilon^2}$. We need to bound the weight of jobs

rejected by $\mathcal{B}$. We can restrict our attention to a fixed machine because $\mathcal{B}$ processes the same set of jobs on a machine as $\mathcal{A}$ does. Further, $\mathcal{A}$ does not reject any job once it gets dispatched to a machine. For rest of the discussion we fix a machine $i^\star$, and bound the weight of jobs which were dispatched by $\mathcal{A}$ to $i^\star$, but got rejected by $\mathcal{B}$. We also fix a density class $d^\star$ and first bound the jobs of density class $d^\star$ which get rejected by $\mathcal{B}$. Let $J^{\mathcal{A}}_{i^\star,d^\star}$ denote the set of jobs of density class $d^\star$ which are dispatched by $\mathcal{A}$ to the machine $i^\star$.

Let $J^{\texttt{rej}}_{i^\star,d^\star}$ denote the subset of jobs in $J^{\mathcal{A}}_{i^\star,d^\star}$ which get rejected by $\mathcal{B}$. First we divide the time line into disjoint intervals $I_1, I_2, \ldots$ with the following properties: (i) for any $j \in J^{\texttt{rej}}_{i^\star,d^\star}$, there is an interval I_r which contains the time period when j was waiting in the queue of machine $i^\star$ (i.e., the time period from r_j to the time when it gets rejected), (ii) for any interval I_r and time $t \in I_r$, there is a job from $J^{\texttt{rej}}_{i^\star,d^\star}$ which is waiting in the queue of machine $i^\star$. We can easily form these intervals by a greedy procedure. For sake of completeness, this procedure is described in Figure 5.

Constructing $I_1, I_2, \ldots$:

1. Initialize t as 0
2. For $r = 1, 2, \ldots$
   (i) Let t_1 be the first time after t when a job from $J^{\texttt{rej}}_{i^\star,d^\star}$ gets dispatched to $i^\star$.
   (ii) Let t_2 be the first time after t_1 when there is no job from $J^{\texttt{rej}}_{i^\star,d^\star}$ in the queue of machine $i^\star$ (in the algorithm $\mathcal{B}$).
   (iii) $I_r \leftarrow [t_1, t_2)$, and update $t \leftarrow t_2$.

Figure 5: Construction of the set of intervals $I_1, I_2, \ldots$.

For an interval I_r, let $J^{\texttt{rej}}_{i^\star,d^\star}(I_r)$ denote the jobs in $J^{\texttt{rej}}_{i^\star,d^\star}$ which are released during I_r. Define $J^{\mathcal{A}}_{i^\star,d^\star}(I_r)$ similarly. We fix an interval I_r and bound the weight of jobs in $J^{\texttt{rej}}_{i^\star,d^\star}(I_r)$. Let $w_{\min}(I_r)$ denote the smallest weight class of a job in $J^{\texttt{rej}}_{i^\star,d^\star}(I_r)$. For a time t and density class d, we define an indicator variable $\mathbf{1}^{\mathcal{A}}_{i,d}(t)$ which is 1 iff $\mathcal{A}$ processes a job of density class d at time t on machine i. Define $\mathbf{1}^{\mathcal{B}}_{i,d}(t)$ similarly. Let $P^{\mathcal{A}}_{i^\star,d^\star}(I_r)$ be the total volume of processing of density class $d^\star$ jobs performed by $\mathcal{A}$ during I_r on machine $i^\star$, i.e.,

$$P^{\mathcal{A}}_{i^\star,d^\star}(I_r) = \int_{t \in I_r} \mathbf{1}^{\mathcal{A}}_{i^\star,d^\star}(t)dt.$$

Claim 6.14. *The total processing size of jobs in $J^{\mathcal{A}}_{i^\star,d^\star}(I_r)$ whose weight class is at least $w_{\min}(I_r)$ is at most $P^{\mathcal{A}}_{i^\star,d^\star}(I_r) + \frac{2\alpha^2 T^\star}{2^{w_{\min}(I_r)}}$.*

Proof. In the schedule $\mathcal{A}$, the jobs in $J^{\mathcal{A}}_{i^\star,d^\star}(I_r)$ will either get processed during I_r or will appear in the queue of machine $i^\star$ at the end of I_r. The total volume of the former quantity is at most $P^{\mathcal{A}}_{i^\star,d^\star}(I_r)$. For the latter quantity, observe that we are interested in weight classes $w_{\min}(I_r)$ and higher. Theorem 6.13 shows that this quantity can be at most

$$\sum_{w \geq w_{\min}(I_r)} \frac{\alpha^2 T^\star}{2^w} \leq \frac{2\alpha^2 T^\star}{2^{w_{\min}(I_r)}}.$$

□

Define another indicator variable $\mathbf{1}_{i^\star,d^\star}(t)$ which is 1 iff both $\mathbf{1}^{\mathcal{A}}_{i^\star,d^\star}(t)$ and $\mathbf{1}^{\mathcal{B}}_{i^\star,d^\star}(t)$ are 1 (the absence of superscript means it is applied to both $\mathcal{A}$ and $\mathcal{B}$). Again, define $P_{i^\star,d^\star}(I_r)$ as the volume of time during I_r for which $\mathbf{1}_{i^\star,d^\star}(t) = 1$. Note a few important points: (i) If $\mathbf{1}_{i^\star,d^\star}(t) = 1$, for some $t \in I_r$, then $\mathcal{B}$ processes a job of density class $d^\star$ and weight class at least $w_{\min}(I_r)$ at time t, because among all jobs of density $d^\star$, it prefers jobs of higher weight, and there is always a job of weight class at least least $w_{\min}(I_r)$ waiting in the queue of machine $i^\star$, (ii) If $\mathbf{1}^{\mathcal{A}}_{i^\star,d^\star}(t) = 1, t \in I_r$ and $\mathbf{1}^{\mathcal{B}}_{i^\star,d^\star}(t) = 0$, then $\mathcal{B}$ processes a job of density class at least $d^\star + \log\left(\frac{1}{\varepsilon}\right)$ at time t.

Now, we want to disregard the part of $P_{i^\star,d^\star}(I_r)$ where $\mathcal{B}$ processes a job which eventually gets rejected. We say that a time t satisfying $\mathbf{1}_{i^\star,d^\star}(t) = 1$ is *bad* if $\mathcal{B}$ processes a job from $J^{\mathtt{rej}}_{i^\star,d^\star}(I_r)$ at time t on $i^\star$. We first show that the total volume of bad time is small.

Lemma 6.15. *The total volume of bad time in I_r is at most $\frac{8\mathtt{length}(I_r)}{\beta}$.*

Proof. For each weight class $w \geq w_{\min}(I_r)$, we bound the volume of bad time at which $\mathcal{B}$ is processing a job of weight class w. Note that for a job of density class $d^\star$ and weight class w, its processing time lies in the range $[2^{w-d^\star}, 2^{w-d^\star-1})$. Let p_w denote $w - d^\star$.

Fix a weight class w. Recall that the algorithm $\mathcal{B}$ divides the time line into segments of length $\beta T^\star/2^w$. Consider such a segment S which intersects with I_r. Let S_L be the segment to the left of S. Any job of weight class w processed by $\mathcal{B}$ during S must have been released in S or S_L (if it were released earlier, $\mathcal{B}$ would have rejected it by the end of S_L). Since $\mathcal{B}$ processes jobs in order of release dates, there will be at most one job j for which $r_j \in S_L$ and $\mathcal{B}$ processes j during a bad time in S. If there were two such jobs j and j' (and say $r_j \leq r_{j'}$), then $\mathcal{B}$ would have completed j before starting j' (note that j can get rejected at the end of S only). But then it could not have rejected j, a contradiction (recall that a job processed during a bad time gets rejected). Hence, the total number volume of bad time in I_r during which $\mathcal{B}$ processes a job of weight class w is at most the number of such segments which intersect I_r times the maximum size of a job of weight class w (and density class $d^\star$), i.e.,

$$\left(\frac{\mathtt{length}(I_r)}{\beta T^\star/2^w} + 2\right) \cdot 2^{p_w} \leq \frac{3 \cdot 2^{p_w} \cdot \mathtt{length}(I_r)}{\beta T^\star/2^w} = \frac{3 \cdot 2^{2w}\mathtt{length}(I_r)}{\beta T^\star 2^{d^\star}},$$

where the first inequality follows from the fact that I_r must contain at least one segment for the weight class $w_{\min}(I_r)$, and so, $\mathtt{length}(I_r) \geq \frac{\beta T^\star}{2^{w_{\min}(I_r)}} \geq \beta T^\star/2^w$.

We sum over all $w \geq w_{\min}(I_r)$, and let $w_{\max}$ be the highest weight class among all jobs of density class $d^\star$. The total volume of bad time in I_r can now be bounded as

$$\sum_{w \leq w_{\max}} \frac{3 \cdot 2^{2w}\mathtt{length}(I_r)}{\beta T^\star 2^{d^\star}} \leq \frac{4 \cdot 2^{2w_{\max}}\mathtt{length}(I_r)}{\beta T^\star 2^{d^\star}} \leq \frac{8 \cdot \mathtt{length}(I_r)}{\beta},$$

where the last inequality follows from the fact that $T^\star \geq 2^{w_{max}} \cdot 2^{w_{\max}-d^\star-1}$ (the weighted size of a job of density class $d^\star$ and weight class $w_{\max}$). ∎

Corollary 6.16. *The total processing time of density class $d^\star$ jobs which are rejected by $\mathcal{B}$ during I_r on machine $i^\star$ is at most*

$$P^{\mathcal{A}}_{i^\star,d^\star}(I_r) + 4 \cdot \varepsilon^2 \mathtt{length}(I_r) - P_{i^\star,d^\star}(I_r)$$

Proof. The result follows from combining Lemma 6.15 and Claim 6.14. Claim 6.14 shows that the total processing time of such jobs is $P^{\mathcal{A}}_{i^\star,d^\star}(I_r) + \frac{2\alpha^2 T^\star}{2^{w_{\min}(I_r)}}$. minus the processing times of such jobs which complete processing in $\mathcal{B}$. Lemma 6.15 shows that the latter quantity is at least $P_{i^\star,d^\star}(I_r) - \frac{8\texttt{length}(I_r)}{\beta}$. Therefore, the processing time of density class $d^\star$ jobs which are rejected by $\mathcal{B}$ during I_r on machine $i^\star$ is at most

$$P^{\mathcal{A}}_{i^\star,d^\star}(I_r) + \frac{2\alpha^2 T^\star}{2^{w_{\min}(I_r)}} + \frac{8\texttt{length}(I_r)}{\beta} - P_{i^\star,d^\star}(I_r).$$

Since $\texttt{length}(I_r) \geq \frac{\beta T^\star}{2^{w_{\min}(I_r)}} = \frac{\alpha^2 T^\star}{\varepsilon^2 2^{w_{\min}(I_r)}}$, the result follows. □

Theorem 6.17. *The total weight of jobs rejected by $\mathcal{B}$ is at most 20ε times the total weight of all jobs.*

Proof. Fix an interval I_r. Corollary 6.16 gives the total weight of density $d^\star$ jobs which get rejected. We write it in a form which will be more useful. For a density class d and time t, let $\mathbf{1}^{\mathcal{B}}_{i^\star,d}(t)$ be the indicator variable which is 1 iff $\mathcal{B}$ processes a job of density class d at time t on machine $i^\star$. Similarly, define $\mathbf{1}^{\mathcal{B}}_{i^\star,\geq d}(t)$ to be 1 iff $\mathcal{B}$ processes a job of density class at least d during t on machine $i^\star$.

We first observe that

$$P^{\mathcal{A}}_{i^\star,d^\star}(I_r) - P_{i^\star,d^\star}(I_r) \leq \int_{t\in I_r} \mathbf{1}^{\mathcal{B}}_{i^\star,\geq d^\star+\log\left(\frac{1}{\varepsilon}\right)}(t)dt,$$

because the LHS is 1 iff at time t, $\mathcal{B}$ processes a job of density class at least $d^\star + \log\left(\frac{1}{\varepsilon}\right)$ at time t. Further,

$$\texttt{length}(I_r) = \int_{t\in I_r} \mathbf{1}^{\mathcal{B}}_{i^\star,\geq d^\star-\log\left(\frac{1}{\varepsilon}\right)}(t)dt,$$

because if at any time $t \in I_r$, there is a job of density class $d^\star$ waiting in the queue at machine i at time t, and so, the processing rule for $\mathcal{B}$ dictates that it cannot process a job of density class less than $d^\star - \log\left(\frac{1}{\varepsilon}\right)$ at time t. Combining the above two inequalities with Corollary 6.16, the total weight of jobs of density $d^\star$ rejected during I_r is at most

$$4\varepsilon^2 \cdot 2^{d^\star+1} \int_{t\in I_r} \mathbf{1}^{\mathcal{B}}_{i^\star,\geq d^\star-\log\left(\frac{1}{\varepsilon}\right)}(t)dt + 2^{d^\star+1} \cdot \int_{t\in I_r} \mathbf{1}^{\mathcal{B}}_{i^\star,\geq d^\star+\log\left(\frac{1}{\varepsilon}\right)}(t)dt.$$

Summing the above for all intervals I_r and noting that these intervals are disjoint for a fixed $d^\star$, the total weight of jobs of density class $d^\star$ rejected is at most

$$4\varepsilon^2 \cdot 2^{d^\star+1} \int_{t} \mathbf{1}^{\mathcal{B}}_{i^\star,\geq d^\star-\log\left(\frac{1}{\varepsilon}\right)}(t)dt + 2^{d^\star+1} \cdot \int_{t} \mathbf{1}^{\mathcal{B}}_{i^\star,\geq d^\star+\log\left(\frac{1}{\varepsilon}\right)}(t)dt.$$

Summing the above for all density classes $d^\star$, we see that the total weight of jobs rejected by $\mathcal{B}$ is at most

$$\begin{aligned}
&4\varepsilon^2 \sum_{d^\star} 2^{d^\star+1} \int_t \sum_{d\geq d^\star-\log\left(\frac{1}{\varepsilon}\right)} \mathbf{1}^{\mathcal{B}}_{i^\star,d}(t)dt + \sum_{d^\star} 2^{d^\star+1} \cdot \int_t \sum_{d\geq d^\star+\log\left(\frac{1}{\varepsilon}\right)} \mathbf{1}^{\mathcal{B}}_{i^\star,d}(t)dt \\
&= 4\varepsilon^2 \int_t \sum_d \sum_{d^\star\leq d+\log\left(\frac{1}{\varepsilon}\right)} 2^{d^\star+1}\mathbf{1}^{\mathcal{B}}_{i^\star,d}(t)dt + \int_t \sum_d \sum_{d^\star\leq d-\log\left(\frac{1}{\varepsilon}\right)} 2^{d^\star+1}\mathbf{1}^{\mathcal{B}}_{i^\star,d}(t)dt \\
&\leq 10\varepsilon \int_t \sum_d 2^d \mathbf{1}^{\mathcal{B}}_{i^\star,d}(t)dt
\end{aligned}$$

Now note that $\sum_i \int_t \sum_d 2^d \mathbf{1}^{\mathcal{B}}_{i,d}(t)dt$ is at most twice the total weight of all jobs processed by $\mathcal{B}$ (a job of class density d has density at most 2^{d+1}). This proves the theorem. □

6.3 Removing the assumption about $T^\star$

So far we have assumed that we know the value of $T^\star$. Now we explain how to get rid of this assumption. Our algorithm starts with an estimate for $T^\star$ and updates it whenever we end up rejecting more than the desired weight of jobs. For a parameter T, let $\mathcal{A}(T)$ denote the algorithm $\mathcal{A}$ when the estimate for $T^\star$ is given by T. Define $\mathcal{B}(T)$ similarly. The modification to the scheduling algorithm is described in Figure 6. Note that when running $\mathcal{A}(T)$ in Step 2(i), the algorithm $\mathcal{A}(T)$ completely ignores the jobs dispatched before j, and so, these jobs do not figure in the calculation of `load`. In fact, these jobs will never get processed by $\mathcal{A}(T)$. We shall refer to each iteration in Step 2 as a *phase*. When we run $\mathcal{B}(T)$ in Step 2(ii), we treat the unfinished jobs of previous phase as being released at the beginning of this phase. So even though such jobs do not affect $\mathcal{A}(T)$ in the current phase, $\mathcal{B}(T)$ may schedule them (or reject them).

Modified Scheduling Algorithm :

1. Initialize $T \leftarrow p_j w_j$, where j is the first released job, and time $t \leftarrow r_j$.
2. **Repeat**
   (i) Run $\mathcal{A}(T)$ on the jobs arriving after j (including j).
   (ii) Run $\mathcal{B}(T)$ alongside $\mathcal{A}(T)$ while treating all unfinished jobs released before j (including j) as being released at time t.
   (iii) If the algorithm $\mathcal{A}(T)$ rejects jobs of total weight more than ε-fraction of the jobs which arrived after j (including j),
   Stop the algorithm $\mathcal{A}(T)$ and update $T \leftarrow 2T$.
   Let j be the last job which was rejected by $\mathcal{A}(T)$. Update $t \leftarrow r_j$.

Figure 6: Scheduling algorithm without any assumption on $T^\star$.

We now analyze the scheduling algorithm.

Theorem 6.18. *The above algorithm is $O(1/\varepsilon^4)$-competitive, and rejects jobs of total weight $O(\varepsilon)$ times the total weight of all jobs.*

Proof. Let $T^\star$ denote the value of the offline optimum. Let T_i be the value of T at the beginning of phase u. First observe that in the last phase $u^\star$, $T_{u^\star} \leq 2T^\star$. Indeed, if T_u becomes larger than $T^\star$, then $\mathcal{A}(T_u)$ will not reject jobs of weight more then ε-fraction of the weight of all the jobs in this phase (the offline optimum for jobs released in this phase can only be at most $T^\star$).

Also, observe that $\mathcal{A}(T_u)$ rejects jobs of weight at most ε-times the weight of all jobs released in this phase, and so, the total weight of jobs rejected by it is within ε-fraction of all the jobs. Now, observe that in a phase u, the total load of jobs of type (w, p) waiting in the queue of a machine i at time t is at most

$$\sum_{u' \leq u} \frac{\alpha^2 T_u}{2^w} \leq \frac{2\alpha^2 T_u}{2^w}.$$

This follows from Theorem 6.13 about the properties of $\mathcal{A}$. Thus, the previous phases worsen the queue size by a factor of 2 only.

Now, we consider $\mathcal{B}$. Suppose it completes a job j in a phase u which was released in phase $u' \leq u$. In a phase u'' between u' and u, j could have waited for at most $\frac{2\alpha^2 T_{u''}}{\varepsilon^2 w_j}$ amount of time (otherwise it would get rejected). So, the total waiting time for this job is at most $\frac{4\alpha^2 T_u}{\varepsilon^2 w_j} \leq \frac{8\alpha^2 T^\star}{\varepsilon^2 w_j}$. Thus, the scheduling algorithm is $O(1/\varepsilon^4)$-competitive.

Further in a phase u, the total weight of jobs rejected by $\mathcal{B}(T_u)$ is at most $O(\varepsilon)$-times the total weight processed by $\mathcal{B}(T_u)$ – this follows from the analysis for algorithm $\mathcal{B}$. The fact that the queue sizes in $\mathcal{A}(T_u)$ are twice the estimate from Theorem 6.13 (because of the effect of previous phases) only doubles the weight of rejected jobs. This proves the theorem. □

7 Extension to `GenWtdMaxFlowTime`

We now extend our result to the `GenWtdMaxFlowTime` problem. Recall that in this problem a a job j has two weights associated with it, the rejection-weight w_j^r and flow-time-weight w_j^f; the first one is used for counting the rejection weight of rejected jobs, while the second one is used in the weighted flow-time expression.

It turns out that almost all the details for the algorithm $\mathcal{A}$ carry over with cosmetic changes in notation to this problem as well, however the algorithm $\mathcal{B}$ needs some change; in particular Lemma 6.15 cannot be applied as it because this is the only place where we critically need the fact that the two weights are same. We now outline the modified algorithm and then the changes that are needed in the analysis. As before we shall assume that the offline optimum value $T^\star$ is known – the details for removing this assumption are exactly as in the case of `WtdMaxFlowTime` problem.

Algorithm $\mathcal{A}$ We define the notion of weight class and density class for each of these two weights. Again, assume wlog that both the weights are powers of 2. We say that a job j is of rejection-weight class w^r if $w_j^r = 2^{w^r}$. Define flow-time-weight class similarly. The rejection-density of a job is defined as w_j^r/p_j; and we say that its rejection-density class is d^r if its rejection-density lies in $[2^{d^r}, 2^{d^r+1})$. Define flow-time-density and flow-time-density class similarly. We now go over the definitions that were used in defining $\mathcal{A}$ and mention the changes in them.

A job j is said to be of type (w^f, d^r) if its rejection-density class is d^r and flow-time-weight class is w^f. For a machine i, time t, and pair (w^f, d^r), let $Q_{i,w^f,d^r}(t)$ denote the jobs of type (w^f, d^r) waiting in the queue of machine i at time t; and define the $\texttt{load}_{i,w^f,d^r}(t)$ as the total weighted remaining processing time of the jobs in $Q_{i,w^f,d^r}(t)$ – if a job j has remaining processing time p'_j, then its *weighted remaining processing time* is defined as $w_j^f p'_j$. When a job j of type (w^f, d^r) arrives at time t, we dispatch it to the machine $i \in S_j$ for which $\texttt{load}_{i,w^f,d^r}(t)$ is minimum, unless for all $i \in S_j$, $\texttt{load}_{i,w^f,d^r}(t) + p_j.2^{w^f} \geq \alpha^2 \cdot T^\star$. If the latter case happens, we reject this job.

It remains to specify which job is processed at any time by a machine. For a time t and machine i, let $(w_i^f(t), d_i^r(t))$ be the pair (w^f, d^r) with the highest $2^{d^r} \cdot \texttt{load}_{i,w^f,d^r}(t)$ value. We process the earliest released job from the queue $Q_{i,w_i^f(t),d_i^r(t)}(t)$ on machine i at time t. Assume a fixed rule of breaking ties.

Algorithm $\mathcal{B}$: Now we describe the modified algorithm $\mathcal{B}$. When a job j arrives at time t, it is dispatched according to $\mathcal{A}$: if $\mathcal{A}$ rejects this job, $\mathcal{B}$ also rejects it; and if $\mathcal{A}$ dispatched it to machine i, then $\mathcal{B}$ also dispatches this job to i. Now, we describe the processing policy for a fixed machine i. Consider a time t. Let $d^{\mathcal{A}}(t)$ denote the rejection-density class of the job processed by $\mathcal{A}$ at time t. Then, $\mathcal{B}$ processes the following job at time t: if there is a job of rejection-density class higher than $d^{\mathcal{A}}(t) + 2\log\left(\frac{1}{\varepsilon}\right)$ in the queue of machine i at time t, then $\mathcal{B}$ processes any such job; otherwise it

processes the job of rejection-density class $d^{\mathcal{A}}(t)$ with the highest flow-time-weight class (it prefers the earliest released job in case of ties). Also note that if the second case happens and there is no job of rejection-density class $d^{\mathcal{A}}(t)$ in the queue of machine i at time t (in $\mathcal{B}$), we can process any job at this time.

The algorithm $\mathcal{B}$ may reject some more jobs. There are two kinds of rejections: (i) immediate rejection: for each triplet (w^f, d^r, p), $\mathcal{B}$ rejects every $\left(\frac{1}{\varepsilon}\right)^{th}$ job of type (w^f, d^r) and size-class p dispatched to it – note that these jobs are rejected as soon as they are released (the algorithm $\mathcal{A}$ is immediate dispatch), (ii) delayed rejection: for every flow-time-weight class w^f, we divide the time line into segments of length $\frac{6(\alpha^2+2)T^\star}{\varepsilon^4 2^{w^f}}$. Suppose a job of flow-time-weight class w^f gets released during such a segment S, and let S' be the segment immediately to the right of S. If the job does not complete processing by the end of S', the algorithm $\mathcal{B}$ rejects the job. This completes the description of $\mathcal{B}$.

Now we outline how the analyses of these two algorithms change.

Analysis of $\mathcal{A}$: The goal is to prove the following extension of Theorem 6.13.

Theorem 7.1. *The algorithm $\mathcal{A}$ rejects jobs of total rejection-weight at most ε times the total rejection-weight of all jobs, and ensures that for any machine i, time t, and pair (w^f, d^r), the total weighted remaining processing time of jobs of type (w^f, d^r) at time t on machine i is at most $\alpha^2 T^\star$. Further, $\mathcal{A}$ is an immediate dispatch algorithm which rejects jobs on arrival only.*

Proof. We give a brief description of the changes that are needed from the proof of Theorem 6.13. A weighted machine interval is again a triplet (I, i, w^f), where w^f corresponds to a flow-time-weight class. Lemma 6.6 holds if we replace weight by flow-time-weight (the lemma gives a lower bound on $T^\star$, which does not depend on rejection at all). Again, we fix a time $t^\star$, and a rejection-density class $d^\star$. The quantities α and γ_l are defined as before. For a machine i, time t, and rejection-density class d^r, $\texttt{load}_{i,d^r}(t)$ is defined as the maximum over all flow-time-weight class w^f, of $\texttt{load}_{i,w^f,d^r}(t)$. The quantities $\texttt{level}(j)$, $J^{(l,d^\star)}$ and $\texttt{maxload}_{i,d^\star}(t)$ are defined analogously.

The weighted machine intervals $\mathcal{I}^{(i,l,d^\star)}$ are constructed as in Figure 4 (in Step 2(iii), w refers to flow-time-weight class). Lemma 6.7 follows without any change (in the proof, we use d to refer to rejection-density class and w to flow-time-weight class). The dual variables are defined as before, and so, Claim 6.8 follows without any change.

For a parameter l and flow-time-weight class w^f, let $\mathbf{I}^{(l,w^f,d^\star)}$ is defined as before, and the quantities $m_{l,w^f,d^\star}, L^{(l,w^f,d^\star)}$ are defined analogously. Inequality (7) holds without any change (if we replace w by w^f). The same applies to Claim 6.10 and inequality (8). The definitions $V_{w,d}$ and $W_{w,d}$ get replaced by (note the subtle change below – the weights in the summation are the flow-time weights in the definition of V but the rejection-weights in the definition of W) :

$$V_{w^f,d^r} = \sum_{j:\texttt{type}(j)=(w^f,d^r)} w^f_j p_j, \quad W_{w^f,d^r} = \sum_{j:\texttt{type}(j)=(w^f,d^r)} w^r_j.$$

$V_{w^f,d^r}(a,b)$ is defined analogously. Claim 6.11 goes through without any changes (if we replace w by w^f and d by d^r). Note that in the proof of this claim, the last line will now read as

$$\frac{2^{d^r}}{2^{w^f}} V_{w^f,d^r} = \sum_{j:\texttt{type}(j)=(w^f,d^r)} \frac{2^{d^r}}{w^f_j} w^f_j p_j = \sum_{j:\texttt{type}(j)=(w^f,d^r)} 2^{d^r} p_j \leq \sum_{j:\texttt{type}(j)=(w^f,d^r)} w^r_j = W_{w^f,d^r}.$$

Proof of Lemma 6.12 now follows without any changes. ∎

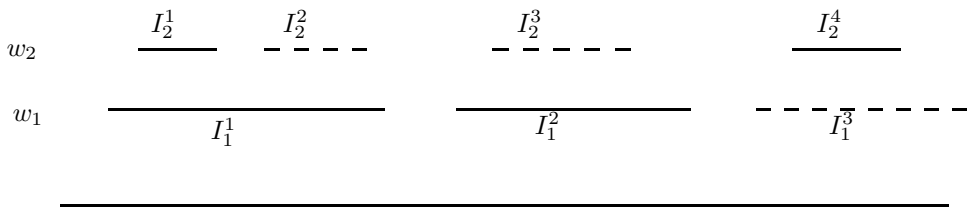

Figure 7: The intervals I_1^1, I_1^2, I_1^3 are maximal intervals where a job of flow-time weight class at least w_1 is waiting in the queue. Out of these $\mathcal{I}(w_1) = \{I_1^1, I_1^2\}$ because these are the intervals where a delayed rejected job of type $(w_1, d^\star)$ is released. Similarly, for a flow-time weight class $w_2, w_2 \geq w_1$, $\mathcal{I}(w_2) = \{I_2^1, I_2^4\}$. Finally, $\mathcal{I} = \{I_1^1, I_1^2, I_2^4.\}$

Analysis of $\mathcal{B}$: We give some intuition about where the analysis differs from that in the case of `WtdMaxFlowTime`. As noted earlier, in the case of `WtdMaxFlowTime`, we needed to bound the volume of bad time slots, i.e., time where $\mathcal{B}$ processes a job, but then decides to reject it later. For this, we crucially needed the fact that the two kind of weights are same. However, we do not have this luxury anymore. This is where we use the immediate rejection – if volume V of jobs (of a certain rejection density class d) are dispatched to a machine i, about εV of this will get immediately rejected. Now, it may happen that $\mathcal{B}$ rejects the remaining jobs after processing them to almost completion, but even then the volume of bad jobs will remain at most $(1-\varepsilon)V$. In the remaining εV time slots (meant for such jobs), $\mathcal{B}$ must process very high density jobs, and so will be able to bound the weight of rejected jobs.

We now proceed with the analysis of $\mathcal{B}$. We only need to bound the total rejection weight of jobs which are rejected. The case of immediate rejection is easy.

Claim 7.2. *The total rejection weight of jobs which are rejected by $\mathcal{B}$ using immediate reject is at most 4ε times the total rejection weight of all jobs.*

Proof. Consider a job of flow-time weight class w^f, rejection density class d^r and size class p. Its rejection weight is at least $2^p \cdot 2^{d^r}$ and at most $2^{p+1} \cdot 2^{d^r+1}$. In other words, any two such jobs will differ in their rejection weight by a factor 4 only. Since $\mathcal{B}$ rejects such jobs after every $1/\varepsilon$ arrivals, the claim follows. □

The case of delayed rejects is more non-trivial. Again, we fix a machine $i^\star$ and rejection density class $d^\star$. For ease of notation, we shall remove the subscripts involving $i^\star$ and $d^\star$. Let J^R denote the set of jobs of rejection density class $d^\star$ which are rejected (using delayed reject) by $\mathcal{B}$. We shall again define a set of disjoint intervals but their definition is more tricky. For a flow-time weight class w^f, we can divide the time line into disjoint set of intervals $\mathcal{I}(w^f)$ such that any interval $I \in \mathcal{I}(w^f)$ satisfies: (i) at any point of time $t \in I$, there is a job of type $(w, d^\star), w \geq w^f$, in the queue of machine $i^\star$ in $\mathcal{B}$, and (ii) at least one job of type $(w^f, d^\star)$ in J^R is released during I. In other words, we first divide the time line into disjoint maximal intervals where we have a job of type $(w, d^\star), w \geq w^f$, in the queue of machine $i^\star$. From these intervals, we pick out those where a job in J^R of type $(w^f, d^\star)$ is released.

Note that if $I \in \mathcal{I}(w_1^f), I' \in \mathcal{I}(w_2^f)$, then either I and I' are disjoint or one is contained in the other. Let $\mathcal{I}$ denote the (containment wise) maximal intervals in $\cup_{w^f} \mathcal{I}(w^f)$, i.e., any two intervals in $\mathcal{I}$ are disjoint and no interval in $\mathcal{I}$ is contained inside any other interval in $\cup_{w^f} \mathcal{I}(w^f)$ (see Figure 7).

Fix an interval $I \in \mathcal{I}$ for rest of the discussion. Let $J^R(I)$ be the set of jobs in J^R which are released during I.

Let $w_{\min}(I)$ be the smallest flow-time weight class of a job in $J^R(I)$ – observe that $I \in \mathcal{I}(w_{\min}(I))$. Further, at every point of time in I, there is a job of type $(w, d^\star), w \geq w_{\min}(I)$ in the queue of machine $i^\star$ in the algorithm $\mathcal{B}$. The maximality of I shows that there is no such job in the queue of machine $i^\star$ just before the left end-point of I.

For a size class p and flow-time weight class w^f, let $P^{\mathcal{A}}_{d^\star,p,w^f}(I)$ denote the total volume of time in I during which $\mathcal{A}$ processes jobs of type $(w^f, d^\star)$ and size class p (on the machine $i^\star$). Let $J^{\mathcal{A}}_{d^\star,p,w^f}(I)$ denote the set of jobs of type $(w^f, d^\star)$ and size-class p which are dispatched by $\mathcal{A}$ to $i^\star$ during I, and $n_{d^\star,p,w^f}$ denotes the cardinality of this set. Note that $\mathcal{A}$ will either process these jobs during I or they will end up in the queue of $i^\star$ at the end of I.

First observe that for any w^f,

$$\sum_p \sum_{j \in J^{\mathcal{A}}_{d^\star,p,w^f}(I)} p_j \leq \sum_p P^{\mathcal{A}}_{d^\star,p,w^f}(I) + \alpha^2 T^\star / 2^{w^f}, \tag{9}$$

because any job in the LHS will either be processed by $\mathcal{A}$ during I or will end up in the queue of type (w^f, d^r) jobs at the end of $i^\star$, and the latter part can have total remaining processing volume at most $\alpha^2 T^\star / 2^{w^f}$ (Theorem 7.1).

Claim 7.3. *The total volume of time in I during which $\mathcal{B}$ processes jobs of type $(w^f, d^\star), w^f \geq w_{\min}(I)$, is at most*

$$(1 - \varepsilon/2)\, P^{\mathcal{A}}_{d^\star}(I) + 2(\alpha^2 + 2) T^\star / 2^{w_{\min}(I)},$$

where $P^{\mathcal{A}}_{d^\star}(I)$ denotes $\sum_p \sum_{w^f \geq w_{\min}(I)} P^{\mathcal{A}}_{d^\star,p,w^f}(I)$, i.e., the total volume of time in I during which $\mathcal{A}$ processes jobs of rejection-density class $d^\star$ and weight-flow-time class $w_{\min}(I)$ or higher.

Proof. Observe that any job of type $(w^f, d^\star)$, $w^f \geq w_{\min}(I)$, processed by $\mathcal{B}$ during I (on machine $i^\star$) has to either: (i) appear in the queue of $\mathcal{B}$ on machine $i^\star$ at the beginning of I, or (ii) dispatched to $i^\star$ by $\mathcal{A}$ during I and not get immediately rejected by $\mathcal{B}$. By definition of I (maximality property), there is no job in (i). We now estimate (ii). Note that for any flow-time weight class w^f and size class p, $\mathcal{B}$ will immediately reject at least $\sum_p (\varepsilon \cdot n_{d^\star,p,w^f} - 1)$ jobs from $J^{\mathcal{A}}_{d^\star,p,w^f}(I)$, and so, the total volume of such jobs rejected by $\mathcal{B}$ is at least

$$\sum_p (\varepsilon \cdot n_{d^\star,p,w^f} - 1) \cdot 2^p \geq \sum_p \varepsilon \cdot n_{d^\star,p,w^f} \cdot 2^p - 2T^\star/2^{w^f} \geq \varepsilon/2 \cdot \sum_p \sum_{j \in J^{\mathcal{A}}_{d^\star,p,w^f}(I)} p_j - 2T^\star/2^{w^f},$$

where the first inequality follows from the fact that the maximum processing time of any job of flow-time-weight class w^f is at most $T^\star/2^{w^f}$. Inequality (9) and the above inequality now imply that the total volume of time during I at which $\mathcal{B}$ processes a job of type $(w^f, d^\star)$ is at most

$$\sum_p \sum_{j \in J^{\mathcal{A}}_{d^\star,p,w^f}(I)} p_j - \varepsilon/2 \cdot \sum_p \sum_{j \in J^{\mathcal{A}}_{d^\star,p,w^f}(I)} p_j + 2T^\star/2^{w^f} \leq (1 - \varepsilon/2) \sum_p P^{\mathcal{A}}_{d^\star,p,w^f}(I) + (\alpha^2 + 2) T^\star / 2^{w^f}$$

Summing over all $w^f \geq w_{\min}(I)$, we get the desired result. □

It follows from the above claim that there must be at least

$$P^{\mathcal{A}}_{d^\star}(I) - \left((1 - \varepsilon/2)\, P^{\mathcal{A}}_{d^\star}(I) + 2(\alpha^2 + 2) T^\star / 2^{w_{\min}(I)}\right) = \varepsilon/2 \cdot P^{\mathcal{A}}_{d^\star}(I) - 2(\alpha^2 + 2) T^\star / 2^{w_{\min}(I)} \tag{10}$$

volume of time during I at which $\mathcal{B}$ is performing jobs of rejection-density class $d^\star + 2\log\left(\frac{1}{\varepsilon}\right)$ or higher – indeed, there is always a job of type $(w^f, d^\star), w^f \geq w_{\min}(I)$ in the queue of $i^\star$ during I, and yet, $\mathcal{B}$ processes such jobs during $(1-\varepsilon/2)\, P^{\mathcal{A}}_{d^\star}(I) + 2(\alpha^2+2)T^\star/2^{w_{\min}(I)}$ out of the possible $P^{\mathcal{A}}_{d^\star}(I)$ volume (during which $\mathcal{A}$ processes such jobs in I). Now, define an indicator variable $\mathbf{1}^{\mathcal{B}}_{i^\star,\geq d^\star}(t)$ which is 1 if $\mathcal{B}$ processes job of rejection-density class at least $d^\star$ at time t on $i^\star$.

Lemma 7.4. *The total processing time of $J^R(I)$, the jobs of rejection-density class $d^\star$ which get delayed rejected by $\mathcal{B}$ during I, is at most*

$$\frac{2}{\varepsilon}\int_{t\in I} \mathbf{1}^{\mathcal{B}}_{i^\star,\geq d^\star+2\log\left(\frac{1}{\varepsilon}\right)} dt + \varepsilon^3 \int_{t\in I} \mathbf{1}^{\mathcal{B}}_{i^\star,\geq d^\star-2\log\left(\frac{1}{\varepsilon}\right)}, \tag{11}$$

Proof. Inequality (10) implies that the first term in the summation above is at least

$$2/\varepsilon\left(\varepsilon/2\cdot P^{\mathcal{A}}_{d^\star}(I) - 2(\alpha^2+2)T^\star/2^{w_{\min}(I)}\right) = P^{\mathcal{A}}_{d^\star}(I) - \frac{4(\alpha^2+2)T^\star}{\varepsilon\cdot 2^{w_{\min}(I)}},$$

and the second term is at least (recall that at all time during I, $\mathcal{B}$ will be processing a job of rejection density class at least $d^\star - 2\log\left(\frac{1}{\varepsilon}\right)$)

$$\varepsilon^3\cdot \texttt{length}(I) \geq \frac{6(\alpha^2+2)T^\star}{\varepsilon\cdot 2^{w_{\min}(I)}},$$

because I contains a job of rejection-weight $w_{\min}(I)$, and so must be as long as one segment corresponding to this rejection-weight class (according to the description of algorithm $\mathcal{B}$). Therefore, the expression in (11) is at least

$$P^{\mathcal{A}}_{d^\star}(I) + \frac{2\alpha^2 T^\star}{\varepsilon\cdot 2^{w_{\min}(I)}},$$

which is at least the total processing size of $J^R(I)$ (using inequality (9) and summing over all $w^f \geq w_{\min}(I)$). □

Rest of the argument follows as in the case of Theorem 6.17 – we sum the expression in (11) over all $I \in \mathcal{I}$, and then over all rejection-density classes $d^\star$ to show that the total rejection weight of rejected jobs is within ε fraction of the total rejection weight of all jobs. Thus, we have shown

Theorem 7.5. *The algorithm $\mathcal{B}$ is $O(1/\varepsilon^6)$-competitive algorithm and rejects jobs of total rejection weight at most $O(\varepsilon)$-times the total rejection-weight of all jobs.*

8 Some Lower Bounds

In this section, we show results on some lower bounds of the `LoadBalancing` and the `MaxFlowTime` problems. The first result shows that for `LoadBalancing`, the trade-off that we obtain between competitive ratio and the fraction of jobs rejected is nearly optimal.

Lemma 8.1. *Given a parameter ε, and an (deterministic) online immediate dispatch algorithm $\mathcal{A}$ for the* `LoadBalancing` *problem, there is an input $\mathcal{I}(\varepsilon)$ consisting of unit size jobs such that $\mathcal{A}$ rejects at most ε-fraction of the jobs and the competitive ratio of $\mathcal{A}$ on $\mathcal{I}(\varepsilon)$ is at least $\Omega\left(\log\left(\frac{1}{\varepsilon}\right)\right)$.*

Proof. The input $\mathcal{I}(\varepsilon)$ has m machines, and will have at most $2m$ jobs. So, the algorithm can reject at most $2\varepsilon m$ jobs. The jobs are released in several phases. At the beginning of phase l, we have a set M_l of machines on which load is at least l. We shall use m_l to denote $|M_l|$.

This is clearly true for $l = 0$ with M_0 being the initial set of machines and $m_0 = m$. Suppose the invariant is true at the beginning of a phase l. During phase l, we shall assume wlog that the algorithm first dispatches *all* the jobs released during this phase, and then rejects some of them. This will not change the competitive ratio of the algorithm because we will only look at the end of a particular phase. We partition the machines in M_l into $\frac{m_l}{2}$ disjoint pairs. For each such pair i, i', we release 2 jobs which can only go on these two machines. Without loss of generality, assume that the algorithm $\mathcal{A}$ dispatches at least one of these two jobs to i. Then we add i to a set X. Thus, X has $\frac{m_l}{2}$ machines. Assuming $m_l \geq 8\varepsilon m$, we can reject at most $2\varepsilon m \leq \frac{|X|}{2}$ jobs dispatched to the machines in X during this phase. Thus, at least half of the machines in X get at least one job during this phase – we let M_{l+1} be these machines. Observe that $\frac{m_l}{4} \leq m_{l+1} \leq \frac{m_l}{2}$. Further the load on a machine in M_{l+1} is at least $l + 1$. Thus, the invariant holds at the end of this phase as well.

Note that we require $m_l \geq 8\varepsilon m$. This will hold if we have only $\left(\log\left(\frac{1}{\varepsilon}\right)\right)$ phases. It is easy to check that the optimal load is at most 2 – during phase l, when we send 2 jobs to a pair of machines i, i', the offline schedule sends these jobs to the machine which does not get added to X. Further, the total number of jobs released is at most $\sum_l m_l \leq 2m$. This completes the proof of the lemma. □

The next result shows that for arbitrary processing times, an immediate dispatch and immediate reject algorithm for the `LoadBalancing` problem will incur high competitive ratio.

Theorem 8.2. *Any online algorithm $\mathcal{A}$ for the* `LoadBalancing` *problem which satisfies immediate dispatch and immediate reject has unbounded competitive ratio, even if it can reject ε-fraction of the jobs.*

Proof. The proof is very similar to that of Lemma 8.1. However, in subsequent phases, jobs sizes will start decreasing, and so, the number of phases will not be bounded by a function of ε. We shall maintain the following invariants for all phases $l = 0, 1, 2, \ldots$:

- At the beginning of phase l, there will be a set M_l of $m_l = \frac{m}{4^l}$ machines such that the load on each of them (in the schedule produced by algorithm $\mathcal{A}$) will be at least $\frac{l}{2}$. Here m is the initial number of machines.
- During phase $l - 1$, we shall release $2^l \cdot m$ jobs each of length $\frac{1}{8^l}$.

These invariants hold at $l = 0$: M_0 is the initial set of machines. Suppose these properties hold for some l. Again, we assume wlog that the algorithm first dispatches all the jobs released in a phase, and then rejects some of them. During phase l, we pair up the machines in M_l into $\frac{m_l}{2}$ pairs. For each such pair of machines i, i', we release $2 \cdot 8^l$ jobs, each of length $\frac{1}{8^l}$, which can only be processed by these two machines. Assume without loss of generality that the algorithm $\mathcal{A}$ dispatches at least 8^l of these jobs to i – we add i to a set X. Hence, X will be a set of $\frac{m_l}{2}$ machines.

Note that the total number of jobs released in phase l is $m_l \cdot 8^l = m \cdot 2^l$. Therefore, the total number jobs released so far is at most $m \cdot 2^{l+1}$, and so, $\mathcal{A}$ can reject at most $\varepsilon \cdot m \cdot 2^{l+1} = 2\varepsilon \cdot m_l 8^l$ jobs released during this phase – note that $\mathcal{A}$ is not allowed to reject jobs released in previous phases during phase l. This implies that at least half of the machines in X will receive at least $\frac{8^l}{2}$ jobs

during this phase – otherwise, the total number of rejected jobs will be at least $\frac{|X| \cdot 8^l}{4} = \frac{1}{8} \cdot m_l \cdot 8^l$, which will be a contradiction (assuming $\varepsilon < 1/16$). Since $|X|/2 \geq m_l/4$, we pick M_{l+1} to be $m_l/4$ such machines. Note that the total load on these machines is at least $\frac{l}{2} + \frac{1}{8^l} \cdot \frac{8^l}{2} = \frac{l+1}{2}$. This proves the invariant for $l+1$.

It is also easy to see that the optimal offline load on any machine is at most 2 (using the same argument as in the proof of Lemma 8.1. Since the number of phases can be arbitrarily large (depending on the value of m), we see that the competitive ratio of $\mathcal{A}$ is unbounded. □

Next we give a strengthening of Lemma 8.1 for the `MaxFlowTime` problem.

Lemma 8.3. *Given a parameter ε, and an (deterministic) online immediate dispatch algorithm $\mathcal{A}$ for the* `MaxFlowTime` *problem, there is an input consisting of unit size jobs such that $\mathcal{A}$ rejects at most ε-fraction of the jobs and the competitive ratio of $\mathcal{A}$ is $\Omega\left(\frac{1}{\varepsilon}\right)$.*

Proof. In our input, we shall release jobs at the end of a time interval I. We assume that the algorithm $\mathcal{A}$ dispatches all the jobs released during I, and rejects jobs only at the end of I. This will be without loss of generality: we will only consider the queue size on a machine at the end of the interval I. If $\mathcal{A}$ was rejecting any job during I, then it could instead dispatch this job arbitrarily and reject it at the end of I. This will not affect the queue length of machines at the end of I.

We shall use the following gadget while building the input.

Claim 8.4. *Suppose at some time t, the algorithm $\mathcal{A}$ is given as input, a set M of $m+1$ machines with i unit sized jobs in the queue of the i^{th} machine, $i = 0, \ldots, m$. Then one can release $m+3$ unit size jobs during the interval $[t, t+2]$ such that at time $t+2$, the load (queue size) on these machines (in some other order) are $0, 1, 2, \ldots, m-1, m+1$.*

Further, there is an offline algorithm which given the machines in M with unit load at time t and the same input as above during $[t, t+2]$ ends up with unit load on all machines at time $t+2$. The maximum load on any machine during this interval is 3.

Proof. We will release several jobs at the same time in a *sequence*, i.e., the next job will be released only after the current one is dispatched by $\mathcal{A}$. We label the machines $M_0, \ldots, M_m$, with M_i having load i at time t in $\mathcal{A}$. All jobs have unit size. The procedure for releasing jobs is described in Figure 8. We first observe that at the beginning of Step 3(i), both M_{i-1} and M_i have load i. Hence, at the end of time t, machine M_i has load $i+1$. Due to Step 4, at the end of time $t+1$, M_m continues to have load $m+1$, but now, M_i has load i for all $i = 0, \ldots, m-1$. Further, no change in load happens due to Step 5, because M_0 already had 0 load. This step is needed to argue about the offline algorithm.

Finally, it is easy to check that there is an offline schedule with the desired properties. In Step 3(ii), the job j_i is dispatched to M_{i-1}. Thus, at the end of time t, M_0 has 3 jobs in its queue, $M_1, \ldots, M_{m-1}$ get 2 jobs, and M_m has one. Step 4 ensures that at the end of time $t+1$, M_0 has 2 jobs in its queue, and the remaining machines have one jobs each in their queue. Finally, in Step 5, M_0 at the end of time $t+2$ also ends up with one job in its queue. This proves the desired result. □

In our input, we shall have Δ machines, where the parameter Δ will be specified later. The jobs are released in $\Delta - 2$ stages. In the beginning of stage l, following invariants are satisfied: (i) In the schedule produced by $\mathcal{A}$, there are l machines with load $1, 2, \ldots, l$, and rest of the machines have 0 load, (ii) In the offline schedule, load is 1 on all machines. Further, the maximum load on any machine till the beginning of stage l was 3.

IncreaseLoad:

1. At time t, release a job j_0 which can only be processed by M_0.
3. **For** $i = 1$ to m **do**
   (i) Release a job j_i at time t which can be processed by $\{M_{i-1}, M_i\}$
   (ii) Assume (upto relabeling) that $\mathcal{A}$ dispatches j_i to M_i.
4. At time $t+1$, release one job j'_m which can only go on machine M_m.
5. At time $t+2$, release one job j''_i for $i = 1, \ldots, m$, such that j''_i can be processed by M_i only.

Figure 8: Dispatching jobs for Claim 8.4

Clearly, this is true at the beginning of stage 0. Suppose this holds true for l – let t_l denote the time at the beginning of stage l. We iteratively call the procedure **IncreaseLoad** defined in Claim 8.4. The procedure is described in Figure 9. Assume that machine M_i has load i for $i = 1, \ldots, l$. The load on other machines remains 0. The procedure uses two other machines, which we call M_0 and M'_0 (note that we have at least $l+2$ machines). During the procedure, we may keep a machine *busy* during a time period. This essentially means that we release one (unit size) at every unit time step during this time period – hence, the load on this machine does not change at all.

ReleaseJobs(Stage l):

1. Initialize time $t \leftarrow t_l$
2. **For** $i = l$ downto 1 **do**
   (i) Machines $M_0, \ldots, M_i$ have load $0, \ldots, i$ (upto reordering) and machines $M_{i+1}, \ldots, M_l$ have load $i+2, \ldots, l+1$ (upto reordering)
   (ii) Call **IncreaseLoad** at time t with machines $M_0, \ldots, M_i$ – these machines get load $0, 1, 2, \ldots, i-1, i+1$ (upto reordering) at time $t+2$
   (iii) **For** $j = i+1$ to l
   Keep machine M_j busy during $[t, t+2]$.
   (iv) $t \leftarrow t+2$.
3. Upto reordering $M_1, \ldots, M_l$ have load $2, \ldots, l+1$, and M_0 has load 0.
   (i) At time t release a job which can be processed by M_0 and M'_0 only, assume $\mathcal{A}$ dispatches this job to M_0 (upto reordering).
   (ii) Keep $M_0, \ldots, M_l$ busy during $[t, t+1]$.

Figure 9: Jobs released in a stage

It is easy to check that at the end of this procedure, we have $l+1$ machines with load $1, 2 \ldots, l+1$ respectively. Suppose there is an offline algorithm in which all machines have load 1 at time t_i. Then Claim 8.4 implies that during Step 2, the maximum load on any machine stays 3, and we end up with 1 load on all machines. In Step 3(i), the job released is sent to M'_0 by the offline schedule, and then during Step 3(ii), M'_0 processes this job to end up with 1 load. This proves that the invariants hold for the next stage as well. Note that the total number of jobs released during this

stage (using Claim 8.4) are at most:

$$1 + \sum_{i=1}^{l} [(i+3) + 2(l-i)] + 1 + (l+1) \leq \frac{3}{2} l^2.$$

Thus, the total number of jobs released till $\Delta - 2$ stages are at most $\frac{\Delta^3}{4}$. At the end of final stage, we have at $\Delta - 1$ machines with load $0, 1, \ldots, \Delta - 1$. Now, the online algorithm can delete at most $\varepsilon \cdot \frac{\Delta^3}{4} \leq \frac{\Delta^2}{8}$ jobs (assuming $\Delta = \frac{1}{2\varepsilon}$). Thus, we will still have at least one machine with total load at least $\frac{\Delta}{2} = \frac{1}{4\varepsilon}$ – indeed, otherwise the total number of deleted jobs will be at least

$$1 + 2 + \cdots + \frac{\Delta - 1}{2} \geq \frac{\Delta^2}{8}.$$

This proves the lemma. Observe that here, the input size is constrained by a function of ε – but we can make it independent of ε by taking multiple copies of the above construction in parallel. □

9 Conclusion and Open Problems

In this paper, we proposed a new model for avoiding the pessimistic bounds arising from competitive analysis of online algorithms for scheduling problems. We could give constant-competitive algorithms for load balancing and minimizing maximum weighted flow-time problems in this model, even though such results cannot be obtained in the speed augmentation model. It is not difficult to show that if there is a single machine, then a policy which rejects every $(\frac{1}{\varepsilon})^{th}$ job (for each weight-class and job size class) and follows HDF rule has competitive ratio within a constant of that when we allow the machine $(1+\varepsilon)$-speed augmentation. Hence, the results which give immediate dispatch algorithms with competitive ratio of $\frac{p}{\varepsilon^{O(1)}}$ for minimizing ℓ_p norm of flow-time in the unrelated machines model with speed augmentation ([5, 28]) can be proved here as well (with an extra $1/\varepsilon$ factor loss in competitive ratio). It is an interesting problem to give algorithms for minimizing ℓ_p norm of flow-time in the rejection model with competitive ratio independent of p. More generally, we feel that more interesting results can be given for other online scheduling problems in this model.

References

[1] Susanne Albers. Better bounds for online scheduling. *SIAM J. Comput.*, 29(2):459–473, 1999.

[2] Susanne Albers and Matthias Hellwig. Semi-online scheduling revisited. *Theor. Comput. Sci.*, 443:1–9, 2012.

[3] Christoph Ambühl and Monaldo Mastrolilli. On-line scheduling to minimize max flow time: an optimal preemptive algorithm. *Oper. Res. Lett.*, 33(6):597–602, 2005.

[4] S. Anand, Karl Bringmann, Tobias Friedrich, Naveen Garg, and Amit Kumar. Minimizing maximum (weighted) flow-time on related and unrelated machines. In *ICALP (1)*, pages 13–24, 2013.

[5] S. Anand, Naveen Garg, and Amit Kumar. Resource augmentation for weighted flow-time explained by dual fitting. In *SODA*, pages 1228–1241, 2012.

[6] James Aspnes, Yossi Azar, Amos Fiat, Serge Plotkin, and Orli Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *J. ACM*, 44(3):486–504, May 1997.

[7] Nir Avrahami and Yossi Azar. Minimizing total flow time and total completion time with immediate dispatching. In *SPAA*, pages 11–18, 2003.

[8] Yossi Azar. On-line load balancing. In *Theoretical Computer Science*, pages 218–225. Springer, 1992.

[9] Yossi Azar, Leah Epstein, and Rob van Stee. Resource augmentation in load balancing. In *SWAT*, pages 189–199. Springer, 2000.

[10] Yossi Azar, Joseph Naor, and Raphael Rom. The competitiveness of on-line assignments. *J. Algorithms*, 18(2):221–237, 1995.

[11] Nikhil Bansal, Avrim Blum, Shuchi Chawla, and Kedar Dhamdhere. Scheduling for flow-time with admission control. In *Proc. ESA, 2003*.

[12] Nikhil Bansal and Kirk Pruhs. Server scheduling in the l_p norm: a rising tide lifts all boat. In *STOC*, pages 242–250, 2003.

[13] Yair Bartal, Stefano Leonardi, Alberto Marchetti-Spaccamela, Jiri Sgall, and Leen Stougie. Multiprocessor scheduling with rejection. *SIAM J. Discrete Math.*, 13(1):64–78, 2000.

[14] Piotr Berman, Moses Charikar, and Marek Karpinski. On-line load balancing for related machines. In *WADS*, volume 1272, 1997.

[15] Niv Buchbinder and Joseph (Seffi) Naor. Fair online load balancing. In *Proceedings of the Eighteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '06, pages 291–298, 2006.

[16] Jivitej S. Chadha, Naveen Garg, Amit Kumar, and V. N. Muralidhara. A competitive algorithm for minimizing weighted flow time on unrelatedmachines with speed augmentation. In *STOC*, pages 679–684, 2009.

[17] Ho-Leung Chan, Sze-Hang Chan, Tak Wah Lam, Lap-Kei Lee, and Jianqiao Zhu. Non-clairvoyant weighted flow time scheduling with rejection penalty. In *SPAA*, pages 246–254, 2012.

[18] Moses Charikar and Samir Khuller. A robust maximum completion time measure for scheduling. In *SODA*, pages 324–333, 2006.

[19] Chandra Chekuri, Ashish Goel, Sanjeev Khanna, and Amit Kumar. Multi-processor scheduling to minimize flow time with epsilon resource augmentation. In *STOC*, pages 363–372, 2004.

[20] Chandra Chekuri, Sungjin Im, and Benjamin Moseley. Online scheduling to minimize maximum response time and maximum delay factor. *Theory of Computing*, 8(1):165–195, 2012.

[21] Chandra Chekuri and Benjamin Moseley. Online scheduling to minimize the maximum delay factor. In *SODA*, pages 1116–1125, 2009.

[22] Leah Epstein and Hanan Zebedat-Haider. Preemptive online scheduling with rejection of unit jobs on two uniformly related machines. *J. Scheduling*, 17(1):87–93, 2014.

[23] Naveen Garg and Amit Kumar. Better algorithms for minimizing average flow-time on related machines. In *ICALP (1)*, pages 181–190, 2006.

[24] Naveen Garg and Amit Kumar. Minimizing average flow-time : Upper and lower bounds. In *FOCS*, pages 603–613, 2007.

[25] Daniel Golovin, Anupam Gupta, Amit Kumar, and Kanat Tangwongsan. All-norms and all-ℓ_p-norms approximation algorithms. In *FSTTCS*, pages 199–210, 2008.

[26] R. L. Graham. Bounds for certain multiprocessing anomalies. *Bell System Technical Journal*, 45:1563–1581, 1966.

[27] Anupam Gupta, Ravishankar Krishnaswamy, Amit Kumar, and Danny Segev. Scheduling with outliers. In *APPROX-RANDOM*, pages 149–162, 2009.

[28] Sungjin Im and Benjamin Moseley. Online scalable algorithm for minimizing ;k-norms of weighted flow time on unrelated machines. In *SODA*, pages 95–108, 2011.

[29] Bala Kalyanasundaram and Kirk Pruhs. Speed is as powerful as clairvoyance. In *FOCS*, pages 214–221, 1995.

[30] Stefano Leonardi and Danny Raz. Approximating total flow time on parallel machines. *J. Comput. Syst. Sci.*, 73(6):875–891, 2007.